Optimal Delegation and Decoupling in the Design of Liability Rules

Ian M. Ayres
Yale Law School

Paul M. Goldbart
University of Illinois at Urbana-Champaign

OPTIMAL DELEGATION AND DECOUPLING IN THE DESIGN OF LIABILITY RULES

Ian Ayres*
Paul M. Goldbart**

INTRODUCTION

Calabresi and Melamed began a scholarly revolution by showing that legal entitlements have two readily distinguishable forms of protection: property rules and liability rules.[1] These two archetypal forms protect an entitlement holder's interest in markedly different ways — via deterrence or compensation. Property rules protect entitlements by trying to deter others from taking. Liability rules, on the other hand, protect entitlements not by deterring but by trying to compensate the victim of nonconsensual takings.[2] Accordingly, the compensatory impetus behind liability rules focuses on the takee's welfare — making sure the sanction is sufficient to compensate the takee. The deterrent impetus behind property rules, however, focuses on the potential taker's welfare — making sure the sanction is sufficient to deter the taker. Thus, disgorgement and prison terms exemplify traditional property rule remedies, while expectation and other compensatory damages fall squarely within the liability rule camp.

Viewing liability rules as a distinct category of entitlement allowed Calabresi and Melamed to identify a missing category in the way courts resolved nuisance disputes. Consider the classic, if somewhat idealized, nuisance dispute between a single "Polluter" and a single "Resident" who is discomforted by the pollution. Prior to *One View of the Cathedral*, courts traditionally chose from among three categories of judgment:

Rule 1: the court issues an injunction against the Polluter;

Rule 2: the court finds the Polluter has created a nuisance but permits pollution to continue provided the Polluter pays damages; or

---

1. Guido Calabresi & Douglas Melamed, *Property Rules, Liability Rules, and Inalienability: One View of the Cathedral*, 85 HARV. L. REV. 1089 (1972).

2. Property rules also often protect entitlements from intentional takings, while liability rules protect them from negligent takings. Hence, you commit a criminal offense if you intentionally break your neighbor's arm (i.e., a property rule protects the integrity of her arm with regard to your intentional taking), but you owe your neighbor compensatory damages if you break her arm through mere negligence (i.e., a liability rule protects the integrity of her arm with regard to your negligent taking).

Rule 3: the court finds the Polluter has not created a nui-
sance and permits pollution to continue without
restriction.

Appreciating the difference between liability and property rules,
Calabresi and Melamed saw that the three approaches above naturally
fit into a two-by-two box. In this stylized box (shown in Figure 1), the
court determines (i) whether the Resident or the Polluter should have
the initial entitlement to control whether pollution will occur,[3] and (ii)
whether to protect this entitlement with a property or a liability rule.[4]

FIGURE 1: CALABRESI AND MELAMED'S TWO-BY-TWO BOX
WITH THE MISSING CATEGORY

		Method of Protection	
		Property Rule	Liability Rule
Initial Entitlement	Resident	Rule 1	Rule 2
	Polluter	Rule 3	[Rule 4]

This traditional tripartite approach to nuisance disputes, however,
leaves one of the boxes empty. Calabresi and Melamed's theory thus
suggests a fourth approach:

---

3. Like Calabresi and Melamed, *see* Calabresi & Melamed, *supra* note 1, at 1091-93, we
use the term "initial" entitlement to convey the idea that, because of the possibility of a sub-
sequent taking, the ultimate entitlement to control whether or not pollution will take place
may belong to a taker.

4. Rule 1, for example, represents a decision in which the court grants the initial enti-
tlement to the Resident and protects it with a property rule. This deters the Polluter, with
the threat of contempt of court, from taking the entitlement nonconsensually. In contrast,
Rule 2 awards the initial entitlement to the Resident but protects it only with a less stringent,
compensation-based liability rule under which the Polluter still might nonconsensually take
the Resident's entitlement but pay compensatory damages for doing so. Rule 2 was famously
deployed in *Boomer v. Atlantic Cement Co.*, 257 N.E.2d 870, 875 (N.Y. 1970). Finally, Rule 3
finds the pollution not to be a nuisance at all, thereby effectively giving the initial entitle-
ment to continue polluting freely to the Polluter under a property rule.

Rule 4:   the court permits pollution to continue unless the
            Resident pays the Polluter damages.[5]

As fate would have it, the fourth box would not stay empty long. In *Spur Industries, Inc. v. Del E. Webb Development Co.*, Justice James D. Cameron of the Arizona Supreme Court — in what has probably become the most cited "coming to the nuisance" court decision — held that the pollution could be enjoined, but only if the developer representing the Residents would "indemnify [the Polluter] for a reasonable amount of the cost of moving or shutting down."[6] Calabresi and Melamed's good fortune in identifying a vacuum that was almost simultaneously filled helped solidify their article's well-deserved fame.

This narrative is well known to the vast majority of tort (and property) professors.[7] However, the legal academy seems much less aware of an analogous phenomenon that has been developing over the last decade. Starting in 1993 with Madeline Morris's article, *The Structure of Entitlements*,[8] a group of about a half dozen scholars began to conceive of "liability rule" protection as a type of "call option."[9] Just as

---

5. In a recent symposium honoring the twenty-fifth anniversary of *One View of the Cathedral*, Symposium, Property Rules, Liability Rules, and Inalienability: *A Twenty-Five Year Retrospective*, 106 YALE L.J. 2083-2213 (1997), Calabresi pointed out that James Atwood had discussed the possibility of a Rule 4-like outcome in a student note that he published in the *Stanford Law Review* three years before Calabresi and Melamed's analysis. *See* Guido Calabresi, *Remarks: The Simple Virtues of the Cathedral*, 106 YALE L.J. 2201, 2204 (1997) (citing James R. Atwood, Note, *An Economic Analysis of Land Use Conflicts*, 21 STAN. L. REV. 293, 315 (1969)).

6. Spur Indus., Inc. v. Del E. Webb Dev. Co., 494 P.2d 700, 708 (Ariz. 1972) (Cameron, V.C.J.). The court's unconditional order seems different than the definition of Rule 4 in the text, which gives the Resident the choice of whether (a) to pay to stop further pollution or (b) not to pay and to allow the Polluter to continue polluting. To harmonize the case with the definition, we must speculate about what would have happened if Del Webb (the developer representing the Residents) had petitioned to void the court order enjoining the pollution as well as the order that Del Webb indemnify Spur (the Polluter). If we believe that the court would have allowed Del Webb to withdraw its initial complaint and thereby void the consequent court orders, it would have in effect given Del Webb the Rule 4 choice — that is, the choice to pay to stop the pollution or not to pay and allow the pollution to continue. At a minimum, future developers will at least realize that suing in this jurisdiction may in effect be choosing to pay for an injunction.

7. Indeed, probably a dozen other articles aver, in one way or another, to this history. *See, e.g.*, Robert C. Ellickson, *Alternatives to Zoning: Covenants, Nuisance Rules and Fines as Land Use Controls*, 40 U. CHI. L. REV. 681 (1973); James E. Krier & Stewart J. Schwab, *Property Rules and Liability Rules: The Cathedral in Another Light*, 70 N.Y.U. L. REV. 440, 442 (1995).

8. Madeline Morris, *The Structure of Entitlements*, 78 CORNELL L. REV. 822, 854-56 (1993).

9. *See, e.g.*, Ian Ayres, *The 1998 Monsanto Lecture: Protecting Property with Puts*, 32 VAL. U. L. REV. 793 (1998); Ian Ayres & J.M. Balkin, *Legal Entitlements as Auctions: Property Rules, Liability Rules, and Beyond*, 106 YALE L.J. 703, 729-33 (1996); Ian Ayres & Eric Talley, *Solomonic Bargaining: Dividing a Legal Entitlement to Facilitate Coasean Trade*, 104 YALE L.J. 1027 (1995); Richard Epstein, *A Clear View of* The Cathedral: *The Dominance of Property Rules*, 106 YALE L.J. 2091, 2093 (1997); Krier & Schwab, *supra* note 7, at 470-75; Saul Levmore, *Unifying Remedies: Property Rules, Liability Rules, and Startling Rules*, 106

the property versus liability rule distinction helped Calabresi and Melamed discover the missing Rule 4, the call-option reconceptualization of the two liability rules (Rules 2 and 4) naturally led to the recognition of two additional "put-option" liability rules.

A call is an option to buy. The option holder can force a sale at the exercise price even if the seller does not want to sell. Rules 2 and 4 have just this type of "call" quality. Under Rule 2, the Polluter has a call option in the form of the choice to pay damages — the choice, in effect, to buy the Resident's right to clean air and, thus, the ability to continue polluting. The Resident possesses the initial entitlement (the right to stop the pollution), but the Polluter can force the Resident to sell this right for a judicially determined amount. Rule 4, on the other hand, reverses the rights of the Polluter and Resident. Under Rule 4, the Polluter possesses the initial entitlement, while the Resident has the call option, and for a judicially determined amount, the Resident can therefore force the Polluter to cede this right to pollute.[10]

Once scholars reconceived traditional liability rules as call options for potential takers, they inevitably began to wonder whether "put-option" rules might not also serve a role. A put is an option to sell. While call options give option holders the choice *to pay* a nonnegotiable amount (the exercise price), put options give option holders the choice *to be paid* a nonnegotiable amount. Call options when exercised give rise to "forced sales," but put options give rise to "forced purchases."

Accordingly, Rule 2 gives the Polluter the choice to pay, and Rule 4 gives the Resident the choice to pay. Scholars considering the classic Rule 2 and Rule 4 court decisions — *Boomer v. Atlantic Cement Co.*[11] and *Spur Industries, Inc. v. Del E. Webb Development Co.*[12] — have often focused only on the difference as to who pays, but a second dimension exists as to who chooses *to force* the payment (i.e., the choice *to be paid*). Rules 2 and 4 concern a choice *to pay*, while the put options concern the choice *to be paid*.

As applied to nuisance disputes, the use of put options suggests two additional rules that give the initial entitlement holder a put option in the form of an option to force a nonconsensual purchase by the other side:

YALE L.J. 2149 (1997); Carol M. Rose, *The Shadow of* The Cathedral, 106 YALE L.J. 2175 (1997) (referring to this feature of liability rules as simply an "option").

10. Property rules (Rules 1 and 3) can also fit easily into a call-option framework: the only difference between a property rule and a liability rule is the difference in the potential taker's exercise price. Property rules set the exercise price so high that no option holder would choose to exercise the option. *See* Ayres & Talley, *supra* note 9, at 1048.

11. 257 N.E.2d 870, 875 (N.Y. 1970).

12. 494 P.2d 700, 708 (Ariz. 1972).

Rule 5: the court permits pollution to continue but also grants the Polluter the option to stop polluting and to receive damages from the Resident; and

Rule 6: the court permits the Resident to enjoin the pollution but also grants the Resident the option to waive this injunctive right in return for damages from the Polluter.

These "put-option" rules — like their "call-option" counterparts — are still "liability rules," designed to compensate the initial entitlement holder for any nonconsensual transfers of that entitlement rather than to deter such transfers altogether. The two differ only in that put options allow the initial entitlement holder to force nonconsensual transfers, whereas under the more traditional call-option liability rule the initially unentitled party can force transfers.

Thus, we can expand Figure 1, as shown in Figure 2, to help flesh out the underlying structure of the six rules:

FIGURE 2: INCORPORATING THE POSSIBILITY
OF "PUT-OPTION" RULES

		Method of Protection		
		Property Rule	Traditional Liability Rule (Initial Entitlement Held Subject to Other Side's Call Option to Pay)	"Put-Option" or "Forced-Purchase" Rule (Initial Entitlement Holder Has Put Option to Be Paid)
Initial Entitlement Held By	Resident	Rule 1	Rule 2	Rule 6
	Polluter	Rule 3	Rule 4	Rule 5

The two new "put-option" rules — Rules 5 and 6 — resemble each other, but under Rule 5, the Polluter has the initial entitlement and put option, whereas under Rule 6, the Resident has the initial entitlement and put option.[13] Thus, the three columns of Figure 2 differ only as to which party has the option to force a nonconsensual transfer of the initial entitlement: under a property rule, neither side has an option to force a transfer; under a traditional liability rule, the unentitled party has the option to force the initial entitlement holder to sell; and under a "put-option" rule, the entitlement holder has the option to force the unentitled party to purchase.

While courts have not used "put-option" rules to resolve nuisance disputes, they have granted put-option awards as a standard response to intentional takings in other areas of tort law.

> Contrary to accepted wisdom, the common law frequently employs "put options" — that is, the right to force a nonconsensual purchase — as a mechanism for protecting entitlements. For example, if Calabresi steals Melamed's watch, Melamed has the option of suing to recover the watch (replevin) or suing to receive the watch's value (trover). Similarly, if Calabresi remains as a holdover tenant in Melamed's apartment, Melamed has the option of suing to enjoin Calabresi's continuing trespass or (at least in some jurisdictions) suing to force Calabresi to pay rent for up to an entire additional year. Finally, if Calabresi builds an encroaching wall on Melamed's land, Melamed has the option of suing to force Calabresi to remove the wall or suing to force Calabresi to permanently buy the land on which he has encroached. In each of these examples, once Calabresi takes Melamed's entitlement, the common law grants Melamed a put option — the option to choose court-determined damages (for permanently ceding the entitlement to the defendant) or injunctive relief (to reacquire the entitlement).[14]

In the nuisance context, instead of giving an Atlantic Cement the (call-option) choice of deciding whether to pay to pollute [Rule 2], a court might give a Boomer the (put-option) choice of whether to be paid (and bear continued pollution) or to refuse payment (and win an injunction against continued pollution) [Rule 6].

Scholars have already made some important progress in analyzing and clarifying the four liability rules (2, 4, 5, and 6). We now know, for example, that under any of four liability rules, courts should set damages equal to the court's best estimate of the value of the party not possessing a call or put option.[15] Authors have also put forward a

---

13. Some authors have reversed the labels of Rules 5 and 6. *See, e.g.,* Ayres, *supra* note 9, at 797 n.9; Levmore, *supra* note 9, at 2163.

14. Ayres, *supra* note 9, at 800.

15. *See* Ayres, *supra* note 9; Louis Kaplow & Steven Shavell, *Property Rules Versus Liability Rules: An Economic Analysis*, 109 HARV. L. REV. 713, 719 (1996).

series of reasons why "put-option" rules might at times dominate "call-option" rules.[16] To date, however, very little of this analysis has addressed a very basic question: to whom should courts award the initial entitlement? Most authors — if they address this question at all — suggest that courts should grant the initial entitlement to the litigant it believes values it more. The *Restatement (Second) of Torts* codifies this standard and directs courts to consider whether "the gravity of the harm outweighs the utility of the actor's conduct" in deciding whether an act constitutes a nuisance.[17] Although the idea of simply giving the entitlement to its higher valuer makes eminent sense in terms of allocative efficiency under the property rule regimes (Rules 1 and 3), it does not make sense when choosing among the four liability rules (Rules 2, 4, 5, and 6).[18] At its core, this Article applies mathematical reasoning to provide practical advice on choosing among liability rules.[19] In doing so, we both simplify and expand the range of judicial choice.

We simplify judicial choice by revealing a deeper structure underlying liability rules. Despite popular belief, maximal allocative efficiency has nothing to do with who is given the initial entitlement. For example, Rule 2 — which confers the initial entitlement on the Resident (subject to the Polluter's call option) — produces the exact ex post allocations, for every possible combination of Polluter and Resident valuations, as Rule 5 — which confers the initial entitlement on the Polluter, along with a put option. Courts concerned about allocative efficiency should therefore focus not on who is given the initial entitlement but rather on delegating the allocative choice to the litigant who is the more efficient chooser.

Our analysis also expands judicial choice by showing that an infinite number of distributive implementations exists within four fundamental classes of liability rules. While any rule within a particular class *allocates the entitlement* between the litigants identically, this multiplicity of implementations within each rule is far from redundant, because each implementation *divides the expected total payoff* differently. Thus,

---

16. *See* Ayres, *supra* note 9, at 801-13; Krier & Schwab, *supra* note 7; Levmore, *supra* note 9, at 2168-72.

17. RESTATEMENT (SECOND) OF TORTS § 826(a) (1979).

18. *See* discussion *infra* Section I.B.

19. For recent discussions on the separate and important question of whether and under what conditions liability rules dominate (or are dominated by) property rules, see Ayres & Talley, *supra* note 9; Ian Ayres & Eric Talley, *Distinguishing Between Consensual and Nonconsensual Advantages of Liability Rules*, 105 YALE L.J. 235 (1995) [hereinafter *Consensual and Nonconsensual Advantages*]; Epstein, *supra* note 9; Kaplow & Shavell, *supra* note 15; and Louis Kaplow & Steven Shavell, *Do Liability Rules Facilitate Bargaining? A Reply to Ayres and Talley*, 105 YALE L.J. 221 (1995) [hereinafter *A Reply to Ayres and Talley*].

In the Conclusion, we return to this issue and show how our analysis also informs our understanding of when property rules will likely dominate liability rules. *See infra* notes 159-164 and accompanying text.

a central finding of this Article is to show that judges can *decouple* allocative and distributive concerns. In other words, a court's decisions about how to maximize ex post allocative efficiency need not affect its decisions about which distribution will best promote equity or ex ante investment.

From the perspective of the initial entitlement holder, liability rules seem to have compensation as its central aim. But from an efficiency perspective liability rules are a means by which an imperfectly informed court can delegate the allocative choice to private litigants who potentially have superior allocative information. For example, a traditional liability rule (Rule 2) can be seen as the court delegating the allocative choice to the Polluter. The Polluter chooses whether to pollute, pay damages to the Resident, and thereby allocate the entitlement to itself, or to allocate the entitlement to the Resident and thus pay nothing. In this Article, we provide a theory of how to delegate such allocative choice optimally.

The four liability rules of Figure 2 (Rules 2, 4, 5, and 6) give rise to just two fundamental allocative classes, corresponding to the identity of the chooser. For convenience, we call these "single-chooser" rules because only one of the litigants determines the entitlement's final allocation. Rules 2 and 5 constitute "defendant-choice" rules because they give the defendant the allocative choice,[20] while Rules 4 and 6 constitute "plaintiff-choice" rules, which give the plaintiff the allocative choice. There is a continuum of defendant-choice rules and a continuum of plaintiff-choice rules, but all the defendant-choice rules (including Rules 2 and 5) allocate the entitlement identically and only differ in the amount of the side payments that are made between the parties. Likewise, there is a continuum of plaintiff-choice rules (including Rules 4 and 6) that identically allocate the entitlement, albeit with different side payments.

We also identify, however, two further fundamental classes of liability rules. We dub these "dual-chooser" rules in contradistinction to the single-chooser rules because both the plaintiff and the defendant potentially impact the final allocation of the entitlement.[21] One class of such dual-chooser rules allots the entitlement to the plaintiff unless

---

20. Under Rule 2 as applied to our Resident and Polluter example, the Polluter has the (call-option) choice either to pollute and pay damages or to refrain and pay nothing; under Rule 5 the Polluter has a (put-option) choice either to refrain from polluting and receive damages or to pollute and receive nothing. Under either rule, the Resident has no power to affect the ultimate allocation except by bargaining with the Polluter in the shadow of these defendant-choice rules.

21. Jack Balkin and Ian Ayres have previously analyzed second-order liability rules, which, like dual-chooser liability rules, give both litigants a voice in the ultimate allocation. *See* Ayres & Balkin, *supra* note 9. But unlike second-order (or higher) rules, which require successively higher exercise prices, the dual-chooser rules at the center of this Article concern a single set of exercise prices on which both litigants base their decisions. We compare and contrast dual-chooser rules and second-order rules more fully in Part III, *infra*.

both parties agree to shift it for a court-determined price to the defendant. In other words, either party has the power to veto allocation of the entitlement to the defendant. Either the plaintiff or the defendant can thereby cause the entitlement to be allocated to the plaintiff. We call this subclass of rules "plaintiff-presumption" allocations. The other subclass of dual-chooser rules allows either party unilaterally to cause the entitlement to be allocated to the defendant, and we correspondingly label these "defendant-presumption" allocations. As with single-chooser rules, we will give multiple option interpretations of these dual-chooser rules and show that there are an infinite number of rules within each of these classes that produce identical entitlement allocations, but which affect how the disputants divide the entitlement's value.

In one sense, the dual-chooser rules represent a kind of centralized planning writ small. Under these rules, the government (that is, a judge) sets a price for the entitlement and then essentially asks both litigants whether they want to trade at that price. While this characterization makes the dual-chooser rules seem esoteric and unworldly, we argue that in some legal contexts, these rules are already being used. Moreover, we show conditions under which a dual-chooser rule systematically produces greater allocative efficiency than either class of single-chooser rules.[22]

The stylized panels of Figure 3 graphically depict our four core allocative equilibria. These panels graph the range of possible plaintiff valuations on the horizontal axis and the range of potential defendant valuations on the vertical axis. If judges possess perfect information about litigants' valuations, they could achieve first-best allocative efficiency by granting entitlements to defendants whenever the defendants' values surpassed the plaintiffs' values. Accordingly, a fully informed judge would apply a decision rule falling along the dashed diagonal lines in the panels.

---

22. Madeleine Morris mentions in passing the theoretical possibility of "dual-chooser," rules, but she does not identify circumstances when they might be advisable nor does she suggest the possibility of their existence in current legal practice. *See* Morris, *supra* note 8, at 854; *see also* Ayres, *supra* note 9, at 823-35 (discussing the possibility of such rules); RONEN AVRAHAM, MODULAR LIABILITY RULES, (John M. Olin Ctr. for Law & Econ., Univ. of Mich., Working Paper No. 01-003, 2001), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=272816 (addressing extensively the rationales for such rules).

Figure 3

(a) Plaintiff-Choice Allocation

v_Δ

0 D_Π v_Π

(b) Defendant-Choice Allocation

v_Δ

D_Δ

0 v_Π

(c) Plaintiff-Presumption Dual-Chooser

v_Δ

$D_{\Pi P}$

0 $D_{\Pi P}$ v_Π

(d) Defendant-Presumption Dual-Chooser

v_Δ

$D_{\Delta P}$

0 $D_{\Delta P}$ v_Π

▬▬▬▬ Entitlement allocated to defendant

▭▭▭▭ Entitlement allocated to plaintiff
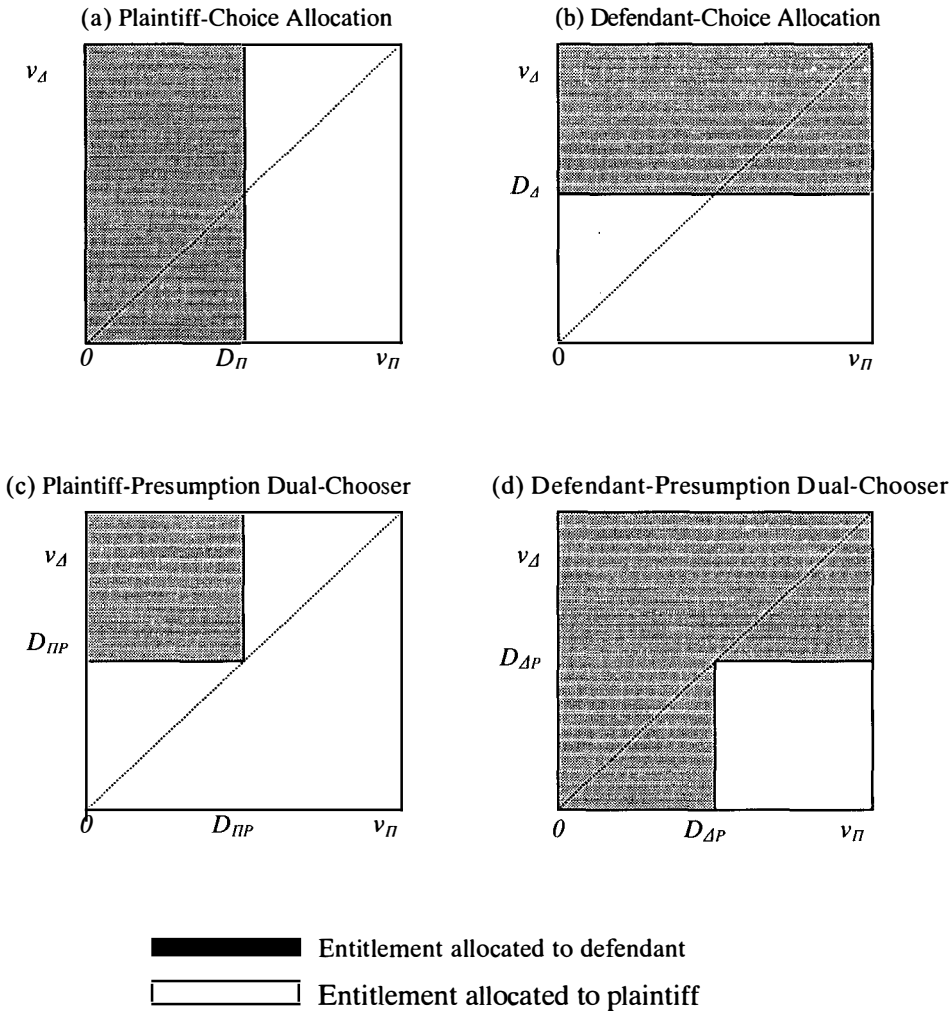
An imperfectly informed judge, however, can instead resort to liability rules to harness the litigants' superior private information regarding their own values.[23] As pictured in panel (a), the class of plaintiff-choice single-chooser allocations vertically partitions the

---

23. *See* Kaplow & Shavell, *supra* note 15, at 725.

valuation space. Consequently, plaintiffs who value the entitlement more than the damage amount will keep the entitlement for themselves under these rules, while plaintiffs who value the entitlement less will allocate the entitlement to the defendants.

Panel (b) depicts the analogous horizontal allocative partition that the defendant-choice single-chooser allocations produce. For this class of rules, the court vests the defendants with the power to allocate the entitlement to themselves when their valuations exceed the damage amount, or to allocate the entitlement to the plaintiffs when their valuations fall short of judicially chosen damages. While these single-chooser allocations do not achieve the first-best allocative efficiency possible under perfect information, they can still enhance allocative efficiency. They give choosers incentives to make the most efficient allocative choices they can, given their private information. In other words, choosers have an advantage over courts because they know their own valuations, even though, like the courts, they know only the probability distributions of their opponents' valuations.

Panels (c) and (d) illustrate the partitions for the dual-chooser rules. Panel (c) shows the plaintiff-presumption dual-chooser allocation, in which the entitlement remains with the plaintiff unless both the plaintiff and the defendant opt to transfer it. Such transfers take place only if the damage amount exceeds the plaintiff's valuation *and* the defendant's valuation exceeds the damage amount. If, on the other hand, the plaintiff's valuation exceeds the damage amount, the plaintiff will choose to keep the entitlement and forego the damage payment the defendant would otherwise have made. Conversely, when the damage amount exceeds the defendant's valuation, the defendant will refuse to accept the entitlement rather than pay the higher damage amount. Similarly, panel (d) shows the analogous defendant-presumption dual-chooser allocation, which ultimately allocates the entitlement to the defendant unless both parties opt for transfer to the plaintiff.

All four classes of rules allocate the upper left quadrant to the defendant and the lower right quadrant to the plaintiff. But there are two crucial differences. First, the four types of rules may differ as to what damage amounts produce optimal efficiency. We will show that the optimal damages for all the rules within a class are the same, but that optimal damages may differ among the classes. One of our tasks in this Article will be to derive the optimal damages for each of the four classes. Second, the rules differ as to how they allocate the "off-diagonal" quadrants (i.e., the lower-left and upper-right quadrants). The dual-chooser allocations give both of these off-diagonal quadrants to the same litigant, while the single-chooser allocations split their ownership, one quadrant to each party. Therefore, to determine which of the four basic liability classes maximizes allocative efficiency in any given situation, a court in effect will need to determine which litigant

will likely make the more efficient allocative choice within each of these off-diagonal quadrants. Accordingly, a second task of the Article is to provide guidance for judges determining which of the four classes in Figure 3 maximizes allocative efficiency, given judges' imperfect information.

The infinite number of distributive implementations within each class allows courts to decouple allocative and distributive decisions, because each implementation within a class divides a constant of expected total payoffs among the litigants differently.[24] Courts choosing from among this expanded set of liability rules should first choose the type of class, that maximizes ex post allocative efficiency, and then choose the implementation within the class that best meets the distributive demands of equity or best aligns ex ante investment incentives.

The core allocative decision involves the optimal number and identity of the choosers. Selection from among these four allocative classes will turn on judicial beliefs about how much the litigants value the entitlement at issue. But instead of being based simply on which litigant has the higher mean value, we show that judges should base this decision more broadly on how they believe these values to be distributed.[25]

We will show that dual-chooser allocations will at times dominate single-chooser allocations and at other times be dominated by them. Dual-chooser allocations produce higher expected ex post efficiency when neither litigant can make a "nuanced" decision alone. Thus, when single-chooser rules might tend to allocate entitlements always to the same litigants (a decision we term "non-nuanced"), dual-chooser rules tend to improve allocative efficiency. Whereas individual litigants, possessing only their own private information, may form an initial opinion about the other litigants' valuations, dual-chooser allocations allow the litigants effectively to exchange information and recognize together instances when they guessed incorrectly. Under these circumstances, two heads can work better than one.

Single-chooser allocations can work more effectively, however, when the court has relatively good information about one litigant's valuation. A court should therefore delegate allocative authority only to litigants who have an informational advantage over the imperfectly informed court. Somewhat counter-intuitively, the litigant who, from the court's perspective, has a less speculative valuation also has less ability to make an efficient allocative choice. A litigant with a less

---

24. In this Article, we generally distinguish between "allocation" of the entitlement at hand and "distribution" of wealth or value.

25. Even after parties have produced evidence of their respective valuations, we imagine that courts will still have residual doubts about their true valuations. Some of these doubts may stem from inevitable uncertainties about the future (e.g., will the factory's new machine work?), while other doubts may simply result from the litigants' strategic incentive to misrepresent their valuations (e.g., "I really value this land because my dog is buried here").

speculative valuation has a smaller informational advantage vis-à-vis both the court and the litigant with damages. Consequently, single-chooser allocations can work more effectively than dual-chooser allocations because litigants who know little more than the court should not wield even partial control over the final destination of the entitlement. When judges can determine one litigant's valuation with relative certainty, they should give the allocative choice only to the other litigant. We will therefore show that the selection between the two single-chooser allocations tends to turn not on which litigant's valuation has a higher mean, but rather on which litigant's valuation has, from the judge's perspective, a higher variance.

Before proceeding, however, we want to emphasize four qualifications to our option approach concerning numerosity, autarky, intentionality, and value correlation. First, we model disputes between only two litigants. Many nuisance examples, however, concern a single source whose pollution adversely affects multiple parties (à la *Boomer* or *Spur*).[26] Nevertheless, scores of nuisance cases centrally revolve around a dispute between only two people.[27] Moreover, we can model even multiple-plaintiff cases in a bilateral fashion if they involve a sufficiently uniform plaintiff class. More fundamentally, our models assume that only litigants contend for the role of highest entitlement valuer. Our option approach powerfully analyzes what economists call "bilateral monopoly,"[28] but in many circumstances third parties will want to enter the notional bidding to possess the entitlement.

Second, our analysis investigates the allocative efficiency of different liability rules in the absence of bargaining (that is, under conditions of "autarky"). This is a restrictive assumption because we should expect that the litigants will bargain in the shadow of the law in many instances. Still, our autarky assumption is supported by recent empiricism suggesting that real world nuisance disputants almost never bargain after courts render judgment.[29] Moreover, as argued by Louis

---

26. Boomer v. Atlantic Cement Co., 257 N.E.2d 870, 875 (N.Y. 1970); Spur Indus., Inc. v. Del E. Webb Dev. Co., 494 P.2d 700, 708 (Ariz. 1972). James Krier and Stewart Schwab provided the best modern discussion of how courts should respond to numerosity problems in choosing among liability rules. *See* Krier & Schwab, *supra* note 7.

27. *See, e.g.*, Copart Indus. v. Consolidated Edison Co., 362 N.E.2d 968 (N.Y. 1977) (pollution from ConEd plant disrupted adjacent new car preparation business); Estancias Dallas Corp. v. Shultz, 500 S.W.2d 217 (Tex. Ct. App. 1973) (air conditioner noise reduced value of single adjacent residence). The class of cases dealing with visual nuisances in particular tends to involve a single plaintiff and a single defendant. *See, e.g.*, Fontainebleau Hotel Corp. v. Forty-Five Twenty-Five, Inc., 114 So.2d 357 (Fla. Ct. App. 1959) (hotel addition obstructed view of adjacent hotel); Amphitheaters, Inc. v. Portland Meadows, 198 P.2d 847 (Or. 1948) (light from night-time dog track interfered with adjacent drive-in movie).

28. Rose characteristically hits the nail on the head when she explains that: "Ayres [is] interested in situations in which two parties are stuck with each other, 'thin' markets instead of 'thick' ones. Neighboring landowners seem to fit that bill." Rose, *supra* note 9, at 2183.

29. *See* Ward Farnsworth, *Do Parties to Nuisance Cases Bargain After Judgment? A Glimpse Inside the Cathedral*, 66 U. CHI. L. REV. 373 (1999).

Kaplow and Steve Shavell, there are reasons to expect that the relative efficiency of different liability regimes when bargaining is not possible will tend to persist when bargaining is possible.[30]

Third, the option model primarily concerns circumstances of intentional taking. While this often (if not usually) describes nuisance disputes, courts often (if not usually) apply liability rules in circumstances of negligent taking. Finally, we assume that the litigants' valuations do not correlate — meaning that an unusually high entitlement valuation for the plaintiff does not entail an unusually high entitlement valuation for the defendant. This often occurs in nuisance suits where the economic benefit to the polluter is orthogonal to the noneconomic detriment to the resident.[31]

Admittedly, these restrictive assumptions limit the applicability of our analysis. But as stressed below,[32] our assumptions fit the stylized facts in a rich set of circumstances. While to keep faith with Calabresi and Melamed we highlight nuisance examples, we believe that contractual disputes provide an even larger class where our analysis obtains. Liability rules (expectation damages) also traditionally protect contractual entitlements; the decision to breach (in other words, to take the promisee's entitlement) is often intentional; and contractual disputes are often unavoidably bilateral.[33] For example, once Epstein agrees to cut Rose's hair, Rose is the only one with whom Epstein can bargain if he wants to buy back his promise.[34] Although our Article does not provide a view of the entire legal "cathedral," the problem of how imperfectly informed courts should resolve bilateral disputes over intentional takings constitutes much more than a theoretical curiosa.

The Article is divided into four parts. Parts I and II analyze single- and dual-chooser allocations, respectively. These Parts derive the op-

---

30. *See A Reply to Ayres and Talley, supra* note 19, at 228-29. The persistence conjecture predicts that a liability rule that has a nonconsensual efficiency advantage under conditions of autarky (i.e., high bargaining costs) will continue under the conditions of low bargaining costs. According to Kaplow and Shavell, headstarts from nonconsensual efficiency will tend to persist when bargaining becomes possible. *Id.* While not denying a persistence effect, however, we believe that different liability rules themselves can induce different types of information disclosure and thereby affect the relative efficiency of rules independent of the autarkic or nonconsensual effect on allocative efficiency. *See Consensual and Nonconsensual Advantages, supra* note 19.

31. With regard to visual nuisances, however, correlated valuations occur more frequently. For example, new structures often derive much of their value from the good views they obstruct. *See, e.g., Fontainebleau Hotel,* 114 So. 2d at 357.

32. *See infra* notes 167-169 and accompanying text.

33. Carol Rose has acutely observed that Ayres's previous option analysis has a contractual shadow as its paradigmatic example: "The contractual relationship has only a discrete number of parties — paradigmatically two — who . . . are stuck with each other." Rose, *supra* note 9, at 2187. Rose is clearly right that contractual renegotiations often involve limited numbers and a bilateral monopoly.

34. *See* Ayres & Balkin, *supra* note 9, at 745. While contract renegotiations may serve as a paradigmatic case of bilateral monopoly, however, other examples exist.

timal damages for the four basic allocations and analyze the conditions under which a particular allocative equilibrium is most likely to be efficient. Part III then relates the dual-chooser allocations to the higher-order allocations previously analyzed by Ayres and Balkin.[35] Finally, Part IV investigates how the preceding results are affected when we relax the most restrictive assumptions of the model. The conclusion shows how our analysis could be applied to improve contract law and also how our analysis illuminates the issue of when liability rules should be used instead of property rules.

I.   SINGLE-CHOOSER RULES

The original pair of liability rules analyzed by Calabresi and Melamed provides examples of the two fundamental single-chooser rules. Rule 2 represents a defendant-choice rule because the defendant (Polluter) chooses who the ultimate entitlement holder will be, deciding whether or not to pay to pollute. Conversely, Rule 4 represents a plaintiff-choice rule because the plaintiff (Resident) chooses who the ultimate entitlement holder will be by deciding whether or not to pay to stop the pollution.

The option recharacterization not only allows us to see the possibility of two more rules, but also clarifies that liability rules (including "put-option" rules) do not differ merely in the way they protect an entitlement. We can also think of them as different ways of dividing the parties' claims to the entitlement. As Carol Rose has written,

> [Under a property rule regime], the entitlement holder has the whole meatball, so to speak, and the other party has nothing — one has property, the other has zip. Under either of the two [call-option] liability rules, on the other hand, the meatball gets split: The factory has an option to pollute (or once exercised, an easement), while the homeowner has a property right subject to an option (or easement). For the sake of simplicity, I will refer to this kind of right as a PRSTO (or PRSTE) for "property right subject to an option (or easement)."[36]

Under this view, options merely segment claims to an entitlement in a different way than, say, a physical or temporal partition.

---

35. *See* Ayres & Balkin, *supra* note 9.

36. Rose, *supra* note 9, at 2178-79. *See generally* Ayres & Talley, *supra* note 9, at 1062-65 (discussing dimensions into which courts might divide an entitlement, including put and call divisions).

Table 1. Possible "Derivative" Divisions of an Entitlement		
Division of Parties' Claims to an Entitlement		
Rule	**Resident's Claim**	**Polluter's Claim**
Rule 1	Entitlement	0
Rule 2	Entitlement – Call	Call
Rule 3	0	Entitlement
Rule 4	Call	Entitlement – Call
Rule 5	– Put	Entitlement + Put
Rule 6	Entitlement + Put	– Put

Table 1 shows how each of the six rules divides the finite claims to an entitlement. Each of the rules maintains conservation of "matter,"[37] in that the sum of the Polluter's and Resident's assets and liabilities under each rule nets out to "Entitlement." As Rose noted, a property rule (such as Rule 1 or 3) gives one party everything and the other person "zip." But other divisions are possible. Rose saw that a "PRSTO" is the same as owning the entitlement subject to a liability (depicted by the "– Call"), in that one might face a sale forced against one's will for a price less than one's actual value.[38] In contrast, the "put-option" rules (Rules 5 and 6) give initial entitlement holders everything that they would have under a property rule *plus* a put option. This necessarily implies that the opposing parties have a liability (denoted by "– Put") — i.e., they might have to purchase against their will for a price that is more than their true value.

---

37. Hohfeld long ago recognized that rights are relational. *See* Wesley Newcomb Hohfeld, *Fundamental Legal Conceptions as Applied in Judicial Reasoning*, 26 YALE L.J. 710 (1917); Wesley Newcomb Hohfeld, *Some Fundamental Legal Conceptions as Applied in Judicial Reasoning*, 23 YALE L.J. 16 (1913). One's right (or entitlement) simultaneously imposes an obligation (or liability) on someone else. The sum of put and call options in any row of Table 1 equals zero, reflecting this Hohfeldian concept.

38. Viewing liability rules as "derivatively" dividing the claims to an entitlement makes explicit the idea that changing the type of "protection" for an entitlement also actually changes the content of the entitlement. *See* Jules L. Coleman & Jody Kraus, *Rethinking the Theory of Legal Rights*, 95 YALE L.J. 1335, 1338-39 (1986). Viewing liability rules as merely asset divisions also calls into question Richard Epstein's claim about the superiority of "property rules." Because property rules also protect both the PRSTO and the call option themselves, *see* Ayres & Balkin, *supra* note 9, at 707, Epstein's claim must relate not just to how to protect entitlements but also to how to bundle collections of strongly protected claims together.

In terms of allocative efficiency, the four liability rules collapse into a choice between just two potential allocations. For any given damage amount and for any given combination of litigant valuations, Rule 2 will allocate the entitlement to the same litigant as Rule 5, and Rule 4 will allocate the entitlement to the same litigant as Rule 6. To see this, consider an entitlement dispute between two litigants — a plaintiff and a defendant. Let v_Π denote the plaintiff's valuation, and let v_Δ denote the defendant's. We assume that both litigants know their own valuations but — like the court — only know the probability distribution of their adversary's valuation. We will refer to the probability density function of the plaintiff's and defendant's valuations, respectively, as $f_\Pi(v)$ and $f_\Delta(v)$, and the distribution means as μ_Π and μ_Δ. Although we will later relax the following assumptions,[39] we also begin by assuming that the litigants' valuations are distributed independently of one another and that the transaction costs (or enmity) are sufficiently high to prevent the litigants from consensually resolving their dispute.

If we let the court-awarded damages equal D, then one can easily see[40] that under Rule 2, the defendant will exercise its call option, pay D damages, and force the plaintiff to sell the entitlement if and only if:

$$v_\Delta > D.$$

Similarly, under Rule 5, the defendant will exercise its put option to force the plaintiff to purchase the entitlement at a price of D damages if and only if:

$$v_\Delta < D.[41]$$

These inequalities indicate that under either of the defendant-choice rules, defendants will allocate the entitlement to themselves when they value it more than the damage amount and will allocate the entitlement to plaintiffs when they value it less. Thus, under the defendant-choice rules, the defendant decides the ultimate allocation of the entitlement independently of the plaintiff's value. Moreover, because the defendant-choice rules produce identical allocations for any given damage amount and combination of litigant valuations, they perforce must also produce identical total payoffs.[42] This means that,

---

39. *See* discussion *infra* Section IV.E.

40. An Appendix available from the authors proves this proposition.

41. If $v_\Delta = D$, the defendant will be indifferent as to whether it exercises either its put or its call options. How the defendant acts on this indifference, however, will not affect the expected total payoffs for each litigant. The inequalities in the text implicitly assume that the defendant's willingness to pay equals its willingness to accept. But see Ayres, *supra* note 9 and *infra* Section IV.E for a discussion of why this assumption might not hold.

42. The payoffs of individual litigants merely equal their valuation of the entitlement should they ultimately receive it, net of any side payments that they make to or receive from the opposing litigants. Because the sum of the side payments must equal zero, the total pay-

given any particular values of D, v_Π, and v_Δ, the defendant-choice rules (Rule 2 and 5) produce equal allocative efficiency.

An analogous result obtains for the plaintiff-choice rules — Rules 4 and 6. Plaintiffs will choose to allocate the entitlement to themselves only when their valuations of the entitlement exceed the damage amount. Like the defendant-choice rules, the allocative efficiency of Rules 4 and 6 remains equal because the physical allocation remains identical for any possible realization of litigants' values.

We can now see that the judicial selection of an initial entitlement holder does not determine allocative efficiency. Rule 2 is equivalent to Rule 5, and Rule 4 is equivalent to Rule 6, even though within each of these pairs the initial entitlement goes to different litigants. Instead, allocative efficiency turns crucially on which litigant has the power to decide the ultimate allocation.

Judicial selection of the initial entitlement holder does, however, seriously affect how the total payoff is distributed between the plaintiff and the defendant.[43] Looking again at Table 1, defendants clearly fare better under the defendant-choice rule that gives them the entitlement plus a put (Rule 5), than under the defendant-choice rule that gives them only a call (Rule 2).[44]

The foregoing suggests that judges choosing among the four liability rules can separate allocative and distributive concerns. Selection of the chooser determines the allocative equilibrium, while selection of the initial entitlement holder determines the distributive equilibrium. To establish a single-chooser regime, judges must (i) select the more allocatively efficient chooser, (ii) select a damage amount, and (iii) select how to distribute the expected total payoff (for example, by choosing between Rules 2 and 5). The remainder of this Section centers on these three tasks.

A. *Selecting Optimal Damages*

While selecting the optimal choosers seems logically prior to selecting the optimal damages, as is often the case in game theory, we can more conveniently solve the model backwards by deriving first the

---

offs for the plaintiff and defendant simply equal the value ascribed to the entitlement by the litigant who ultimately owns it.

43. Note the parallel to the Coase Theorem. In a world without transaction costs, the judicial assignment of entitlements does not affect the allocative equilibrium but does affect the distribution of payoffs. Likewise, an autarkic model of liability rules, where high transaction costs make bargaining prohibitive, can produce a similar result. Ayres and Balkin demonstrated that liability rules form a kind of auction mechanism that would mimic the results of Coasean negotiations. *See* Ayres & Balkin, *supra* note 9, at 727-29.

44. For example, if the plaintiffs and defendants valuations are independently and uniformly distributed between $0 and $100, and if damages are set equal to $50 (which we later show is the optimal amount), then Rule 5 will yield an expected defendant payoff of $62.50, while Rule 2 will yield the defendant only $12.50. *See* Ayres, *supra* note 9, at 806.

optimal damage amount for a particular chooser and then identifying the more efficient chooser. Derivation of optimal damages is conceptually straightforward. Simply setting damages equal to the nonchooser's expected value maximizes the allocative efficiency of single-chooser rules.[45] This means that the (allocatively) optimal damage amount for the defendant-choice allocations will equal:

$$D_\Delta = \mu_\Pi,$$

and the damage amounts for the plaintiff-choice allocations will equal:

$$D_\Pi = \mu_\Delta.$$

These results apply quite generally and hold true regardless of the distribution of litigants' valuations, as long as the litigants' distributions are independent.[46] Courts may still have legitimate concerns about equity and other aspects of efficiency, such as inducing adequate ex ante incentives to create the entitlement in question. But we will show that courts have a separate ability to establish almost any desired distributive equilibrium, and this strongly counsels in favor of damages that will maximize the size of the pie.

Damages that equal nonchoosers' expected valuations maximize the allocative efficiency of a single-chooser allocation because they give the choosers an incentive to allocate entitlements to themselves only when their valuations surpass the mean valuations of the nonchoosers. Similarly, they give choosers an incentive to allocate entitlements to nonchoosers only when their own valuations lie below the mean valuations of the nonchoosers. Damages equivalent to nonchoosers' mean valuations thus harness the private information of choosers (who are reasonably assumed to know more about their own valuations), and on average induces an efficient ultimate asset allocation.[47] For example, if $50 represents the plaintiff's mean valuation,

---

45. *See* Kaplow & Shavell, *supra* note 15 (deriving this result with regard to call-option Rules 2 and 4).

46. If the distributions do correlate, courts should still set the optimal single-chooser rule damages at the nonchooser's mean value — but because this mean value now occurs as a function of the chooser's value, optimal damages will occur at the fixed point at which damages equal both the nonchooser's value and the chooser's value. *See* discussion *infra* Section IV.F.

47. As discussed in the next section, Section I.C, *infra*, the single-chooser rules give the nonchoosers a fixed expected payoff whether or not the choosers exercise their options. For example, under Rule 2, the plaintiffs as nonchoosers receive an expected value equal to their mean valuations if defendants do not exercise their call options because plaintiffs will retain the entitlements. On the other hand, plaintiffs will receive their expected valuations in cash if defendants do exercise their call options and pay damages set equal to plaintiffs' expected valuations. Fixing the nonchoosers' expected payoffs makes choosers the residual claimants of gains arising from exercise of their options. The nonchoosers resemble bondholders in corporate law who are fixed claimants on the assets of the firm, while choosers resemble the shareholders who hold the residual claims. Because choosers internalize all the marginal effects of their choice, single-chooser rules give them, like the shareholders in corporate law, incentives to choose the allocation that maximizes allocative efficiency, conditional on their superior private knowledge.

damages of $50 under either Rules 2 or 5 will, on average, encourage the defendant to take the entitlement only when such an allocation is efficient.[48] Even when judges do not know as much as individual litigants, single-chooser liability rules allow judges to exploit the private information of one of the litigants.

B.   *Selecting the Chooser*

The second judicial task that affects allocative efficiency involves selection of the chooser. Although the put and call implementations of a given single-chooser rule produce identical allocations, nothing guarantees that a plaintiff-chooser rule will produce the same allocation as a defendant-chooser rule for any given damage amount and combination of litigant valuations. In this Section, we provide some guidance as to which party is the more efficient chooser — i.e., which chooser will produce the highest total expected payoff.[49]

As a first step to identifying the more efficient chooser, courts must calculate the total, combined plaintiff and defendant expected payoff for a particular single-chooser allocation (assuming optimal damages). We can express these payoffs analytically as:

$$E[\text{Total Payoff}] = \mu_{nc} + \int_{\mu_{nc}}^{\infty}(q_c - \mu_{nc})f_c(q_c)dq_c$$

where the subscripts c and nc refer to the chooser and nonchooser, respectively. We derive the foregoing expression in an appendix,[50] but we can also express it in more stylized terms as:

Expected Total Payoff = nonchooser's mean + chooser's call.   (A)

Intuitively, the chooser has the power to allocate the entitlement to the nonchooser (in which case the total payoff will equal the nonchooser's) or the chooser can exercise its option of allocating the entitlement to itself (in which case the total payoff will equal the nonchooser's mean plus the expected amount by which the chooser's mean exceeds the damage amount (equivalent to the nonchooser's mean)). This option (which we label "chooser's call") of allocating the entitlement to itself when the chooser's privately known value exceeds

---

48. Defendants under Rule 2 will choose to pollute and pay $50 only if their valuations exceed the plaintiffs' mean valuations; defendants under Rule 5 will forego payments of $50 only if their valuations exceed that amount.

49. *See* Krier & Schwab, *supra* note 7 (suggesting that the smaller, less numerous party makes the more efficient chooser). The concept of "efficient chooser" is an autarkic analog to the concept of the more "efficient briber," discussed originally by Calabresi & Hischoff. Guido Calabresi & John T. Hirschoff, *Toward a Test for Strict Liability in Torts*, 81 YALE L.J. 1055, 1060 (1972).

50. Appendix, *supra* note 40.

the nonchooser's mean is the reason that liability rules systematically produce higher efficiency than property rules, in the absence of bargaining.

But it turns out that the expected total payoff can also be expressed as:

Expected Total Payoff = chooser's mean + chooser's put.   (B)

Thus, we can also think of the single-chooser allocation as allowing the chooser to keep the entitlement (yielding a payoff on average equal to the chooser's mean) or to put the entitlement to the nonchooser when the chooser's privately known value falls below the nonchooser's mean.

Finance cognoscenti will recognize the equivalence of these two payoff expressions (A and B) as an implication of the "put-call parity formula," which teaches that whenever one can express a value in terms of an implicit call, one can also express the same value in terms of an implicit put.[51]

When applied to both plaintiff- and defendant-choice allocations, these two expressions of the expected total payoff give rise to four equivalent comparisons for determining which litigant is the better chooser. The plaintiff-choice rules — Rules 4 and 6 — will operate more efficiently than the defendant-choice rules — Rules 2 and 5 — if and only if the following equivalent conditions hold:

$$\text{(i) } \operatorname{call}_\Pi(\mu_\Delta) > \operatorname{put}_\Delta(\mu_\Pi);$$

$$\text{(ii) } \operatorname{put}_\Pi(\mu_\Delta) > \operatorname{call}_\Delta(\mu_\Pi);$$

$$\text{(iii) } \mu_\Pi - \mu_\Delta > \operatorname{put}_\Delta(\mu_\Pi) - \operatorname{put}_\Pi(\mu_\Delta); \text{ and}$$

$$\text{(iv) } \operatorname{call}_\Pi(\mu_\Delta) - \operatorname{call}_\Delta(\mu_\Pi) > \mu_\Pi - \mu_\Delta;$$

where, for example, "$\operatorname{call}_\Pi(\mu_\Delta)$" represents the expected value of a call option for a plaintiff when the exercise price of the call (reported in the parentheses) is equal to the defendant's mean valuation.[52] Two points immediately merit emphasis. First, these four inequalities merely restate one another so that probability distributions that satisfy one of the inequalities must automatically satisfy the others. Second,

---

51. *See* RICHARD A. BREALEY & STEWART C. MYERS, PRINCIPLES OF CORPORATE FINANCE 557-59 (5th ed. 1996) (defining put-call parity); Ayres & Talley, *supra* note 9, at 1047 (applying put-call parity to liability rules). The put-call parity formula is often stated in the following form:

Call + Exercise Price = Put + Underlying Asset.

As applied to single-chooser rules, since the exercise price of the options is set equal to the nonchooser's mean value, the left-hand side of this equation equals the first expression (equation *A*) for the expected total payoff. And since the expected value of the underlying asset to the chooser is its mean, the right-hand side of this equation equals the second expression (equation *B*) for the expected total payoff.

52. Analogously, "$\operatorname{put}_\Pi(\mu_\Delta)$" represents the expected value of a put option for a defendant when we set the exercise price of the put equal to the plaintiff's mean valuation.

the exercise price of the options is always set at the mean valuation of non-option holder (i.e. the nonchooser). This latter point is just an implication of the last section's finding that optimal damages should be set at the nonchooser's mean value.

The most important implication of these comparisons, which becomes particularly evident in the first two inequalities, is that the litigant who from the court's perspective has the more speculative[53] valuation will tend to make the more efficient chooser. Again, this will not surprise finance cognoscenti. A central result of derivative theory is that options tend to be more valuable as the value of the underlying asset becomes more volatile. The litigant that the court perceives to have the more speculative value is the litigant with the greater informational advantage.

Even though we assume the litigants' values are privately known, those values can be uncertain from the court's perspective. And the court in delegating the allocative choice should grant the allocational option to the litigant with the greater informational advantage. The litigant whose value is more difficult for the court to estimate is likely to be the more efficient option holder because this litigant in choosing whether to exercise her options indicates whether her value is likely to be larger than the other litigant. Even though it is strange to think of the option holders with privately certain but socially uncertain values, from the perspective of a court (or social planner) with imperfect information the social value of the option is more naturally analogous to the value of financial option — equaling (as shown above) to the likelihood that the underlying asset (the litigant's value) will exceed the exercise price (damages).

Indeed, when the court must select between a litigant with a commonly known valuation and a litigant with a valuation known to fall only within a particular probability distribution, the litigant with the known value never constitutes the more efficient chooser.[54] This makes intuitive sense. Litigants with commonly known valuations have no additional information to bring to the decision. If courts gave such litigants the power to allocate the entitlement (and set damages equal to the more speculative mean valuation of the opposing litigant), litigants with known valuations would therefore always allocate enti-

---

53. We use the word "speculative" in a special sense. Our model assumes that each litigant places a non-speculative, certain dollar value on its ownership of the entitlement while the court and the adversaries must speculate (via a probability distribution) about how much the entitlement is valued by the litigant. But in many real-world contexts, however, the litigants themselves may not know precisely their own valuations. A polluter, for example, may also have to speculate about how much profit a polluting factory will yield in the future. We will later relax this assumption that litigants are perfectly informed about their own valuation, *see infra* Section IV.E, but for now it is sufficient to say that our core qualitative results go through. The more efficient chooser will continue to be the litigant with the larger informational advantage.

54. *See* Appendix, *supra* note 40.

tlements to themselves when their own valuations exceed their opponents' mean valuations, or would allocate entitlements to their opponents when their known valuations fall below their opponents' mean valuations. Thus, a court that grants allocative power to the litigant with the known, nonspeculative valuation also effectively selects the ultimate entitlement holder. If the other litigant's valuation happens to differ from the mean of the probability distribution, however, the entitlement could end up in the wrong hands as a consequence. In contrast, courts that grant allocative power to litigants with the speculative valuations, and set damages equal to the valuations of the nonspeculative litigants, will always achieve first-best allocative efficiency. Because speculative litigants know when their own valuations deviate from the mean of their probability distributions, they can prevent what might otherwise become a misallocation of the ultimate entitlement. In other words, speculative litigants will allocate entitlements to themselves when they know, privately, that their valuations exceed those of the nonspeculative litigants, and will also know to allocate entitlements to the nonspeculative litigants when their own valuations are lower.

In inspecting the latter two inequalities, (iii) and (iv), listed above, one might initially believe that the relative mean valuations of the litigants should also bear on the selection of the more efficient chooser. But this turns out not to be true. First, as an algebraic matter, notice that in the third inequality, a higher mean plaintiff valuation, relative to the defendant's mean, seems to *increase* the likelihood that the plaintiff will choose more efficiently. Yet, in the fourth inequality, a higher plaintiff mean seems to *decrease* the likelihood that the plaintiff will choose more efficiently. Both cannot prove true, because these inequalities are mathematically equivalent.[55] As it turns out, neither intuition proves true because the values of the put options and call options also change in ways that offset the direct impact of any change in the litigants' relative mean valuations. For example, if the plaintiff's mean rises relative to the defendant's mean (holding constant the variance of both litigants' valuations), both the plaintiff's call and the defendant's put will become more valuable. This shift in relative mean valuations drives both these options further "into the money."[56]

---

55. *See id.*

56. From a court's perspective, a plaintiff whose mean valuation exceeds the defendant's mean valuation will more likely exercise a call option because more of the plaintiff's probability distribution will exceed the call's exercise price, which equals the defendant's mean. Similarly, the defendant will more likely exercise a put option because more of the defendant's probability distribution will fall below the put's exercise price, which equals the plaintiff's mean. In other words, the value of call or put options depends in part on the odds that the litigants will actually exercise them. The odds that litigants will exercise their options in turn lie in direct proportion to the odds that their own mean valuations fall above or below their opponents' mean valuations.

As a general matter, option value arises as a function of both a "volatility" effect and an "in the money" effect. As a result of the "volatility" effect, options generally increase in value when an option holder's valuation increases in volatility. Under the "in the money" effect, options generally increase in value when (a) a call option holder's valuation exceeds, or (b) a put option holder's valuation falls short of, the option's exercise price. In working through the four inequalities, however, one can see that the "in the money" effect on option value affects both sides of each inequality in the same manner. An increase in the plaintiff's mean valuation, relative to the defendant's valuation, increases both sides of inequality (i) and decreases both sides of inequality (ii). In inequality (i), for example, the "in the money" effect on the plaintiff's call exactly matches its effect on the defendant's put.[57] Consequently, a shift in relative mean valuations should not — given constant volatility — push a court toward either favoring or disfavoring a particular litigant as the more efficient chooser.

Courts, in selecting the more efficient chooser, should focus on the second moment of the distribution (the variance), not the first moment (the mean) and select the litigant who, from the court's perspective, has a higher variance of valuation. This finding suggests that the *Restatement (Second) of Torts* is misguided in suggesting that courts should assign the initial entitlement in nuisance disputes to the party with the higher valuation. The *Restatement*'s allocation principle makes eminent sense where the court merely chooses between the two property rules, Rule 1 and 3. The principle fails to maximize allocative efficiency, however, when the court contemplates a choice between the call-option versions of the liability rules (Rules 2 and 4). Although we demonstrated above that selection of the initial entitlement holder need not affect the ultimate allocation, this holds true only in a world that employs both put-option and call-option liability rules — so that the choice of the initial entitlement holder can be decoupled from the choice of the more efficient chooser. But the world contemplated by the *Restatement* does not decouple these two choices. In the *Restatement*'s world (where the only liability rules that judges award are the call-option versions — Rules 2 and 4), a rule that the litigant with the higher perceived value should be the initial entitlement holder is tantamount to a rule that the litigant with the lower perceived value should be the chooser — via call-option Rule 2 or 4. Our model does not support the *Restatement*'s position. Rather, we propose that when there are two litigants, the more efficient chooser will tend to be the

---

57. These effects come into play in inequalities (iii) and (iv) as well. A rise in the plaintiff's mean relative to the defendant's mean also increases both sides of inequalities (iii) and (iv). In inequality (iii), for example, the left-hand side $(\mu_{\Pi} - \mu_{\Delta})$ increases, but the right-hand side also increases as the shift in mean causes "$put_{\Delta}(\mu_{\Pi})$" to increase and "$put_{\Pi}(\mu_{\Delta})$" to decrease.

chooser who from the court's perspective has the greater variance of potential private valuation.

From a pragmatic perspective, our two single-chooser allocations work together so that the court can avoid estimating speculative damages. We have suggested that courts should select as chooser the litigant with the more speculative value and then set damages equal to the less speculative mean of the nonchooser.[58]

In a less reductive model, the court's selection of the more efficient chooser will turn on more than just the relative volatility of valuations. For example, the litigants' relative levels of risk aversion might affect a court's selection. The variance of the choosers' expected payoffs will tend to be lower than the nonchoosers' because choosers have greater control over their destiny.[59] Thus, efficiency-minded courts, other things held equal, might consider higher risk aversion, as well as more speculative damages, as desirable qualities for a potential chooser.

At other times, the more efficient chooser will be the party that is less numerous.[60] If a single source of pollution harms multiple residents, it may be difficult to implement a plaintiff-choice allocation. We will return to this issue below when we provide a more extended discussion of how numerosity affects our analysis.[61]

C. *Selecting the Distribution*

The first two judicial tasks — selecting the chooser and selecting the damages — are sufficient to determine the allocative equilibrium for any particular realization of litigant valuations. That is to say, these two judicial decisions establish who will ultimately posses the entitlement. These decisions do not, however, determine the distributive equilibrium. For example, a court may finally decide on a defendant-choice rule, but it still needs to select between Rules 2 and 5 in deciding how to divide the payoff pie between the litigants. In this Section, we show that a court can divide the expected total payoff as it likes between the two litigants — unfettered by considerations about alloca-

---

58. We agree with Kaplow and Shavell that the speculative nature of damages should not exclude them from the court's assessment, *see* Kaplow & Shavell, *supra* note 15, at 728-32, but our approach simplifies the judicial task by requiring only an assessment of the less speculative litigant's average valuation.

59. *See* Ayres, *supra* note 9, at 812-13, 821 & n.81. For example, if the litigants' values are both uniformly distributed between $0 and $100, then a single-chooser rule will give the chooser a fixed amount 50% of the time and a $50 range of values the other 50% of the time, but the rule will give the nonchooser a fixed payoff 50% of the time and a $100 dollar range of values the other 50% of the time. We derive this more formally in the Appendix, *supra* note 40.

60. *See* Krier & Schwab, *supra* note 7, at 471-72.

61. *See* discussion *infra* Section IV.D.

tive efficiency. In other words, the court can decouple questions of distribution from questions of allocation.

Obviously, a chooser given only a call option to buy an entitlement will receive a lower payoff than a chooser given an entitlement already plus a put option. Accordingly, the put implementations of the single-chooser allocations yield a relatively larger payoff for choosers and a relatively smaller payoff for nonchoosers than the call implementations of the single-chooser allocations.

But we can go even further in characterizing the relative payoffs of each litigant. Interestingly, the single-chooser allocations grant the nonchooser an expected payoff that is independent of the entitlement's ultimate allocation.[62] The nonchooser's expected payoff does not vary with respect to the chooser's actual allocative choice. More specifically:

1) Put implementations of the single-chooser allocations fix the nonchoosers' expected payoffs at zero whether or not choosers exercise their put options; and

2) Call implementations of the single-chooser allocations, on the other hand, fix the nonchoosers' expected payoffs at their mean valuations, μ_{NC}, whether or not choosers exercise their call options.

To see this, first consider the put implementation of the single-chooser allocations (where the chooser has both the entitlement and a put option). If the chooser does not exercise the put, the nonchooser's payoff is zero with certainty because the nonchooser neither receives the entitlement nor make a payment. Instead, the nonchooser remains at status quo. Alternatively, if the chooser does exercise the put, the nonchooser's expected payoff is zero on average because the nonchoosers must pay μ_{NC} for something that they value on average at μ_{NC}.

---

62. In analyzing the distributive consequences of different rules, one must distinguish between two different informational perspectives. The ex ante perspective assesses the litigants' expected payoffs before they acquire information about their own valuation of the entitlement. The *interim* perspective assesses the litigants' expected payoffs after they identify their own valuations but before learning the other side's valuation. From the ex ante perspective, the possibility of facing the other side's put option does not produce a negative expected payoff. Under Rule 6, a Resident may force a Polluter to purchase the right to pollute, but if the court sets the price at the Polluter's expected valuation, the Polluter's expected payoff equals zero. In contrast, from the *interim* perspective, the possibility of facing to the other side's put option can produce a negative expected payoff if, for example, under Rule 6 a Polluter happens to have a lower than average valuation. The ex ante perspective also matches the perspective of the court in deciding how to divide the expected payoffs between sets of observationally equivalent litigants. This perspective may at times also bear relevance to the formation of efficient investment incentives, if the parties don't know their actualized valuation until after the court creates the entitlement. Nevertheless, the *interim* perspective will potentially drive both investment and allocation decisions in other contexts.

Similar reasoning shows that call implementations fix the nonchoosers' expected payoff at μ_{NC}.[63]

Thus, call implementations of single-chooser allocations give the nonchooser an expected payoff that is μ_{NC} higher than put implementations yield. And since we know that put and call implementations of a single-chooser allocation produce identical expected *total* payoffs, equal to $\mu_{NC} + call_C$, it must be true that:

1) The chooser's expected payoff in put implementations equals the total expected payoff, $\mu_{NC} + call_C$; and

2) The chooser's expected payoff in call implementations equals the value of a chooser's call, $call_C$.

The varying distributions of the put and the call implementations give policymakers more flexibility in pursuing nonallocative goals; but, if only two implementations existed, courts would still possess only a rather crudely hewn distributive instrument. Distributively minded lawmakers might at times prefer to sacrifice ex post allocative efficiency in favor a more preferred distribution by selecting an inefficient chooser.[64]

Fortunately, courts can radically expand the class of single-chooser allocations by "convexifying" the payoffs between the call and the put implementations. We refer to this as our *convexity* result. It turns out that the two single-chooser allocations can be implemented by not just four rules but by a double continuum of implementations that give policymakers the power of smoothly distributing the expected joint surplus between the litigants as they see fit.[65]

The key here is to distinguish between the damages that choosers must pay if they allocate entitlements to themselves, which we denote as D_C, from the damages that choosers must pay if they allocate enti-

---

63. If a chooser exercises the call option, the nonchooser receives μ_{NC} with certainty. If a chooser does not exercise the call option, the nonchooser simply retains the entitlement, which the nonchooser values on average at μ_{NC}.

Of course, once nonchoosers come to know their private value of the entitlement, they will not be indifferent about whether the chooser exercises its option or not. But the court (and other policy makers) will never be privy to this privately held information, so the best that a distributively concerned court can do is to adopt the type of implementation that tends, on average, to produce a preferred distribution.

64. For example, imagine that the plaintiff is the less efficient chooser. Then the court, in making the defendant the chooser, could give the plaintiff an expected payoff of zero (through a defendant put implementation — Rule 5) or μ_{Π} (through a defendant call implementation — Rule 4). But if the court felt that the plaintiff deserved an expected payoff between zero and μ_{Π}, it might decide to make the plaintiff the chooser via a call implementation — so that the plaintiff's expected payoff would fall in this intermediate level ($0 < call_{\Pi} < \mu_{\Pi}$).

65. Other authors have suggested a proliferation of liability rules on different grounds, *see, e.g.*, Levmore, *supra* note 9, at 2171 (showing how liability rules could require different levels of negligence before taking), but, with our convexity result, we are the first to focus on distributive flexibility.

tlements to nonchoosers, which we denote as D_{NC}. As shown in Table 2, a call implementation of a single-chooser allocation sets D_C equal to the nonchoosers' mean, μ_{NC}, and sets D_{NC} equal to zero. This means that the choosers must pay to take the entitlement nonconsensually but pay nothing if they leave the entitlement in the hands of the nonchoosers. In contrast, a put implementation sets D_C equal to zero and sets D_{NC} equal to $-\mu_{NC}$, which means that the choosers pay nothing to retain the entitlement but receive a payment of μ_{NC} if they put the entitlement into the hands of the nonchoosers.

We normally think of an option as requiring a monetary payment only if the option holder exercises the option. But when the puts and calls are described in terms of the amounts paid by the option holder contingent on *either* choice, we open a range of new single-chooser implementations that still induce the choosers to make identical allocations. For example, as shown in Table 2, merely averaging the damages in the put and the call implementations creates a new implementation. This new implementation requires the choosers to pay half the nonchoosers' mean value when taking the entitlement but entitles the choosers to a payment of half the nonchoosers' mean value if they allocate the entitlement to the nonchoosers. We refer to this implementation as a "pay-or-be-paid" rule. This third rule induces identical allocations as the put and the call implementations. For example, instead of asking the plaintiff Residents (under *Spur*-like application of Rule 4) whether they are willing to pay one million dollars to enjoin the defendants' pollution, a "pay-or-be-paid" rule might ask the plaintiff to elect whether she prefers to pay half a million dollars to stop pollution or receive half a million dollars from defendant to accept continued pollution. In either case, the choosers will only allocate the entitlement to themselves if their private value is greater than the nonchoosers' mean value because self-allocation requires the choosers to incur a direct cost of half the nonchoosers' mean value and an opportunity cost of half the nonchoosers' value, for a sum total in this case of a million dollars.[66] Indeed, any implementation with damages such that:

$$D_C - D_{NC} = \mu_{NC},$$

where the damages choosers pay for taking the entitlement less the damages choosers receive for handing over the entitlement equals the mean of the nonchoosers, will produce an identical allocation.

---

66. Opportunity cost corresponds to "the price that you pay for things that you might have done." Ian Ayres, *Analyzing Stock Lock-Ups: Do Target Treasury Sales Foreclose or Facilitate Takeover Auctions?*, 90 COLUM. L. REV. 682, 688 n.19 (1990) (quoting BILLY JOEL, *Only the Good Die Young*, *on* THE STRANGER (Columbia Records 1977)).

Table 2. Expanding the Class of Single-Chooser Implementations				
	Damages		**Expected Payoff**	
Implementation	D_C	D_{NC}	**Chooser**	**Nonchooser**
Call	μ_{NC}	0	$call_C$	μ_{NC}
Put	0	$-\mu_{NC}$	$call_C + \mu_{NC}$	0
"Pay-or-Be-Paid"	$\mu_{NC}/2$	$-\mu_{NC}/2$	$call_C + (\mu_{NC}/2)$	$\mu_{NC}/2$
"Pay-or-Pay"	$\mu_{NC} + call_C$	$call_C$	0	$call_C + \mu_{NC}$
Alpha (for all $0 \le \alpha \le 1$)	$\alpha(\mu_{NC} + call_C)$	$\alpha(\mu_{NC} + call_C) - \mu_{NC}$	$(1-\alpha)(\mu_{NC} + call_C)$	$\alpha(\mu_{NC} + call_C)$

One can also create implementations that produce even more lop-sided distributions than either the put or the call implementations without undermining the choosers' incentive to make the efficient choice. For example, Table 1 shows that when the damages choosers pay to take the entitlement, D_C, equals the nonchoosers' mean valuation plus a call price, $\mu_{NC} + call_C$, and the payment that the choosers receive for handing over the entitlement, D_{NC}, equals the call price, $call_C$, a court can restrict the choosers to a zero expected payoff. The court can thus give the entire expected total payoff to the nonchoosers.[67] We call this a "pay-or-pay" rule because the choosers must pay the non-choosers regardless of whether they allocate the entitlement to themselves or to the nonchoosers; allocating to the nonchoosers merely reduces the amount that the choosers must pay. For example, rather than ask Polluters under Rule 2 whether they will pay one million dollars for the right to pollute, a pay-or-pay rule might ask them to choose between paying $1.2 million for the right to pollute or paying $0.2 million to forego polluting. As with the foregoing pay-or-be-paid rule, choosers (absent wealth and framing effects)[68] will under either rule allocate the entitlement to themselves only if their private valuations exceed the nonchoosers' mean valuations.

---

67. To verify this result, just consider the nonchooser's expected payoff. If the chooser keeps the entitlement, the nonchooser receives $\mu_{NC} + call_C$ with certainty; and if the chooser allocates the entitlement to the nonchooser, the nonchooser receives $call_C$ and retains the entitlement, which on average it values at μ_{NC}.

68. *See* discussion *infra* Section IV.E.

The put and the pay-or-pay implementations are of particular interest because they give these alternatives (while maintaining the single-chooser rule's allocative efficiency) that allow judges to allocate all of the expected value to either litigant. The put implementation gives the entire expected payoff to the chooser, while the pay-or-pay gives the entire expected payoff to the nonchooser.[69] In the absence of bargaining, the property rules, Rules 1 and 3, also allocate entire payoffs to one litigant or the other,[70] but the distributively extreme liability rules, the put and pay-or-pay implementations, systematically dominate their property rule counterparts because they produce larger expected total payoffs.

By creating convex combinations of the two different types of damages under these two distributively extreme implementations of the single-chooser allocations, courts gain the ability to adopt any intermediate division of the expected total payoff. As shown in the row labeled "alpha" in Table 2, setting

$$D_C = \alpha \left(\mu_{NC} + call_C \right), \text{ and}$$
$$D_{NC} = \alpha \left(\mu_{NC} + call_C \right) - \mu_{NC},$$

the court effectively distributes a fraction α of the expected total payoff to the nonchooser and the remaining fraction, $1 - \alpha$, to the chooser. Thus, by choosing the number α between zero and one, this "alpha" implementation gives the court unfettered flexibility in dividing the pie between the litigants, without undermining the chooser's incentives to maximize the pie's size. For example, when α equals 0.9, the court can distribute ninety percent of the expected total payoff to the nonchooser and only ten percent to the chooser while still encouraging the chooser to allocate the entitlement to the litigant with the higher valuation. This alpha class includes both pay-or-be-paid rules (in which α takes on low values), and pay-or-pay rules (in which α takes on sufficiently higher values).

The reader will no doubt have realized that these new-fangled pay-or-be-paid and pay-or-pay distributive rules are identical to rules in which the court simply implements a traditional call-option rule — but in addition requires a lump-sum payment between the litigants regardless of what allocation the chooser made. Under this reconceptualization, the court would set two prices — an allocative price (D_A), which would determine the amount the chooser would pay if it allocated the entitlement to itself, and a distributive price (D_D), which would determine the amount the chooser would pay (or be paid if D_D was negative) regardless of its allocative choice. While these alternative implementations are mathematically equivalent, we conjecture

---

69. Courts can effect even more lopsided distributions by shifting the D_C and D_{NC} damages amounts up above the pay-or-pay or down below the put amounts.

70. *See* Ayres & Talley, *supra* note 9, at 1062-65; Rose, *supra* note 9, at 2178-79.

that bundling the allocative and distribute prices together into discrete rules such as the pay-or-be-paid rule will be less jarring to judges than an explicitly separate distributive lump sum amount.[71] Indeed, even readers who recognized this lump sum addition to the traditional call option liability rule probably did not recognize that the put-option rules are mathematically nothing more than call-option rules in which courts set D_D equal to $-\mu_{NC}$.[72] The very novelty of the put-option implementations suggests that the foregoing analysis has not always been obvious to legal scholars. Indeed, one of the payoffs of this Article is to show that put-option rules are not esoterica but merely implementations of call-option rules with a lump-sum side payment.

In the past, courts may have been tempted to accommodate distributive considerations by simply raising the call-option exercise price of the traditional liability rule. Courts could accomplish this, for example, by requiring an Atlantic Cement to pay more than the Residents' expected valuation to acquire a pollution right.[73] The foregoing analysis shows that such an award would needlessly sacrifice allocative efficiency by inducing Atlantic Cement to take too infrequently on the altar of equity. In this example, a better way to increase the Residents' payoff would be to institute a pay-or-pay rule — which would simultaneously increase the amount that Atlantic Cement would pay even if it chose not to pollute (D_{NC}) and increase the price it must pay if it chose to pollute (D_C). More generally, we have shown that it is misguided for courts to deviate from setting damages at the nonchooser's mean valuation as a way to improve distributive equity.

None of these implementations — including the infinity of alpha implementations — affect the variance of expected payoff for either the chooser or the nonchooser. Accordingly, the adoption of any one

___

71. There is a sense, of course, in which the court in setting both D_C and D_{NC} is setting two prices. Other law and economic scholars have considered decoupling the amount the defendant pays from the amount the plaintiff receives. *See* A. Mitchell Polinsky and Yeon-Koo Che, *Decoupling Liability: Optimal Incentives for Care and Litigation*, 22 RAND J. ECON. 562 (1991). The type of decoupling this Article proposes, however, differs because the side payments from the litigants always net to zero.

72. Rule 6, for example, uses a put-option rule that gives the plaintiff both the initial entitlement and the option to put the entitlement to the defendant at a price equal to defendant's mean valuation. But the same allocation and distribution could be implemented by a *Spur*-like Rule 4 (which gives the initial entitlement to the defendant but gives the plaintiff a call option to buy at the defendant's mean valuation) if in addition the court also ordered the defendant to pay her mean valuation to the plaintiff regardless of the plaintiff's choice (i.e., $D_D = -\mu_{NC}$). Under the call implementation with the lump-sum mean side payment, plaintiffs would end up with (i) the entitlement and no net dollar transfer if they allocate the entitlement to themselves, or (ii) a net dollar transfer equaling the defendant's mean value, if plaintiffs allocate the entitlement to the defendant. This, of course, yields the same outcome as under the put implementation.

73. *Cf.* Boomer v. Atlantic Cement Co., 257 N.E.2d 870, 875 (N.Y. 1970). At some point, however, an additional increase in the exercise price of the defendant's call would not increase the plaintiff's expected payoff because the defendant would never agree to pay such a price.

of the distributive implementations within a particular single-chooser allocation would not affect a court's selection of a chooser on the basis of relative risk aversion. A court's selection of the more efficient chooser — whether it be motivated by considerations of numerosity, variance, or risk aversion — is independent of the court's distributive choice.

This convexity result allows lawmakers to decouple distributive concerns from ex post allocative efficiency concerns and thus frees courts to divide the expected total payoff in order to further either equity or to generate better ex ante investment incentives. From an equity perspective, courts might, for example, fine-tune the extent to which it rewards a plaintiff's coming-to-a-nuisance claim with an expected payoff. The option perspective also provides courts with an amount to distribute that exceeds even the mean valuation of the higher valuing party. Single-chooser allocations produce an expected total payoff that systematically exceeds the expected valuation of either party, and a court concerned about matters of equity can now make a more explicit decision about which litigant has a stronger equitable claim to this option value.

From an efficiency perspective, a court's division of the expected total payoff can exert important effects on the parties' ex ante investment incentives. If lawmakers place a high value on the Residents' incentives to invest in their land, a grant of the entire payoff in the form of an initial entitlement plus a put option will give them a much stronger incentive to invest than merely giving them a call.[74] Thus, expanding the class of plaintiff-chooser and defendant-chooser allocations not only allows lawmakers to better accommodate competing equity and efficiency concerns, but can also help lawmakers better accommodate competing efficiency concerns — that is, the concern with ex ante investment efficiency and the concern with ex post allocative efficiency.

While we have shown in our simple model that the courts distributive choice need not affect the chooser's allocative decision making, there is a sense in which the distributive rules may affect the nonchooser's decision making. In some litigation contexts, the nonchooser may take an initial action that triggers the chooser's subsequent (put or call) option. For example, the next section will show that a common law put rule can be seen as a two-stage process: in the first stage, the defendant intentionally takes, thereby signaling a willingness to face the plaintiff's put option in the second stage. Such two-stage liability regimes exemplify what we call dual-chooser rules because both liti-

---

74. *See* Ayres, *supra* note 9, at 807; Ayres & Balkin, *supra* note 9, at 733. The ex ante investment effects of traditional liability rules are explored by Lucian Arye Bebchuk, *Property Rights and Liability Rules: The Ex Ante View of the Cathedral*, 100 MICH. L. REV. (forthcoming Dec. 2001).

gants have an impact on the ultimate allocation. The goal of the next section is to analyze what damages maximize the allocative efficiency of such rules, and to assess the conditions under which dual-chooser allocations dominate (or are dominated by) single-chooser allocations.

II. DUAL-CHOOSER RULES

As mentioned in the introduction, there are two or more fundamental allocations that judges might choose in resolving disputes. We refer to the legal regimes that give rise to these allocations as "dual-chooser" rules as opposed to single-chooser rules because both litigants have a potential impact on how the entitlement is allocated. As with the single-chooser rules, dual-chooser rules give rise to two different allocations. In what we call the "plaintiff-presumption dual-chooser" allocation, either litigant can unilaterally cause the entitlement to be allocated to the plaintiff. If either the plaintiff or the defendant prefers that the entitlement be allocated to the plaintiff (along with any court-determined side-payment), then the plaintiff receives the entitlement. The entitlement is allocated to the defendant only if both litigants agree to a court-crafted trade. But because either litigant is sufficient to ensure a plaintiff allocation, the plaintiff-presumption equilibrium disproportionately favors the entitlement to the plaintiff. This pro-plaintiff allocation bias appears in panel (c) of Figure 3, in which three out of four quadrants are allocated to the plaintiff.

Conversely, the defendant-presumption dual-chooser rule allocates the entitlement to the defendant unless both plaintiff and defendant choose to allocate the good to the plaintiff (along with the associated court-determined side-payment). Because either litigant can veto a plaintiff allocation, this type of rule disproportionately favors allocation of the entitlement to the defendant, as shown in panel (d) of Figure 3.

All four of these basic allocations — the two single-chooser allocations and the two dual-chooser allocations — arise from what we label "single-price" allocations. In these allocations, the court establishes both a single allocative price[75] and the identity of the chooser or choosers, and then the court sits back and lets the chooser(s) allocate the entitlement.

A number of different option interpretations give rise to these dual-chooser implementations. Here we examine two plaintiff-presumption dual-chooser implementations, but transposing the iden-

---

75. In one sense, of course, the court sets two prices in setting both D_C and D_{NC}. From an option perspective, however, the court sets only one relevant allocative price, the difference between these amounts ($D_C - D_{NC}$). As discussed above, the court can implement the decoupling of allocative and distributive concerns alternatively by explicitly setting a single allocative price, D_A, and a single distributive price, D_D. *See supra* notes 71-72 and accompanying text.

tity of the plaintiff and defendants can create analogous defendant-presumption implementations. Consider the following two divisions of entitlements:

$$\Pi: \quad E + \text{put}(D_{\Pi P}) - \text{put}(D_{\Pi P}) \qquad \Delta: \quad -\text{put}(D_{\Pi P}) + \text{put}(D_{\Pi P}) \qquad \text{(A)}$$

$$\Pi: \quad \text{call}(D_{\Pi P}) - \text{put}(D_{\Pi P}) \qquad \Delta: \quad E - \text{call}(D_{\Pi P}) + \text{put}(D_{\Pi P}) \qquad \text{(B)}$$

In the first, "double-put" regime, (A), the court gives the plaintiff the initial entitlement plus a put option to receive a price of $\$D_{\Pi P}$, but subject to the defendant's put option to sell the entitlement back to plaintiff, also for $\$D_{\Pi P}$. This allocations allows either litigant to veto the transfer of the entitlement to the defendant. The plaintiff can veto the transfer of the entitlement by refusing to exercise its initial put, but even if the plaintiff exercises this put, the defendant can veto the transfer by putting the entitlement back to the plaintiff. We can easily derive the equilibrium choices, or strategies, that this allocation induces. The plaintiff will put the entitlement only if:

$$v_\Pi < D_{\Pi P},$$

and the defendant will put the entitlement back only if:

$$v_\Delta < D_{\Pi P.}$$

where v denotes the precise, privately known valuation of each litigant, as opposed to their respective mean speculative valuations, μ. As illustrated in panel (c) of Figure 3, these two conditions together mean that in equilibrium, the defendant will receive the entitlement if and only if:

$$v_\Pi < D_{\Pi P} < v_\Delta.[76]$$

Alternatively, under the second regime, (B), the court gives the plaintiff a call option to buy the defendant's initial entitlement but subject to the defendant's put option to sell. This allocation allows *either* litigant to force a transfer of the entitlement to the plaintiff with a payment of $\$D_{\Pi P}$ from plaintiff to defendant. As with regime (A), the defendant will retain the entitlement if and only if:

$$v_\Pi < D_{\Pi P} < v_\Delta.$$

This assures that the allocations of these two rules will be identical for any possible combination of plaintiff and defendant private valuations. As with the single-chooser allocations, however, these different plaintiff-presumption dual-chooser implementations produce very dif-

---

76. The court could implement another option scheme to allow either party to veto a transfer of the entitlement to the defendant by:

$$\Pi: E - \text{call}(D_{\Pi P}) + \text{call}(D_{\Pi P}) \qquad \Delta: + \text{call}(D_{\Pi P}) - \text{call}(D_{\Pi P}).$$

In the text's double-put rule, the plaintiff can in a sense offer to sell the entitlement for $\$D_{\Pi P}$ by initially putting, and the defendant can accept this offer by choosing not to put the entitlement back. In contrast, this footnote's double-call rule allows the defendant to offer to buy the entitlement by exercising its initial call option, and the plaintiff can accept this offer by choosing not to call the entitlement back.

ferent distributions. By inspecting the initial entitlement division, it is apparent that the plaintiff's expected payoff increases if the plaintiff receives the initial entitlement (and potentially receives subsequent payment for transfer of the entitlement) than if the plaintiff does not receive the initial entitlement (and can only gain it by paying a court-determined amount).[77]

At first blush, these dual-chooser allocations seem esoteric and otherworldly, but as we intimated at the end of the last section, several common law settings already provide at least the potential for joint-chooser allocations. For example, if Laurel intentionally encroaches on Hardy's land, Hardy traditionally has (in addition to damages for the prior encroachment) the choice of ejectment (an injunction forcing Laurel to remove the encroachment) or permanent damages (forcing Laurel to permanently purchase[78] the encroaching property at a court-determined amount).[79] Instead of beginning the analysis with Hardy's allocative choice (of taking back the land or forcing a sale to Laurel), it may be more useful — at least with regard to deliberate encroachment — to think of the law as implementing a dual-chooser allocation. From this perspective, the encroacher signals a willingness to buy property at a court-determined price by intentionally encroaching on it, and the plaintiff indicates a willingness to sell at this same court-determined price by suing for damages instead of ejectment. More generally, in property law, a temporary taking or impairment of another's right often has the effect of triggering a put option.[80] When a deliberate taking, with the prospect of paying compensatory damages, results in a plaintiff's choice between compensation and restoration of the entitlement, these dual choices of the litigants implement a plain-

---

77. The option interpretation of these dual-chooser rules makes clear that they are constrained versions of what one of us has previously referred to as "second-order rules." *See* Ayres & Balkin, *supra* note 9. The dual-chooser rules constrain these second-order rules by setting the exercise price of both options to an identical amount, D. Part III, *infra*, of this Article analyzes and compares the unconstrained analog to this constrained, dual-chooser version.

78. The court may thereby force an encroacher not only to compensate the landowner for the temporary, past encroachment, but also to purchase a "permanent" right to encroach on the land in the future, even if the encroacher would actually prefer to remove the encroachment instead.

79. *See* Pile v. Pedrick, 31 A. 646, 647 (Pa. 1895) (granting plaintiff the choice of damages for "permanent trespass" or an injunction to remove the "offending ends of the stones," after finding that defendant's brick wall encroached one and three-eighths inches underneath the plaintiff's property).

80. Accordingly, we can apply a similar two-stage analysis to (i) a tenant's choice to hold over, followed by a landlord's choice to force an additional term lease, *see* JESSE DUKEMINIER & JAMES E. KRIER, PROPERTY 431 (3d. ed. 1993); RESTATEMENT (SECOND) OF PROPERTY: LANDLORD AND TENANT § 14.4 cmt. f (1977), or (ii) an intentional taking of chattel followed by the owner's choice of trover (i.e., compensatory damages) or replevin (i.e., an injunction ordering the chattel's return), *see* DUKEMINIER & KRIER, *supra*, at 105-06.

tiff-presumption dual-chooser allocation — in that the entitlement will be allocated to the defendant only if both the defendant's private valuation surpasses, and the plaintiff's private valuation falls below, the expected court-awarded damages.

Indeed, under current law, courts may have difficulty implementing true plaintiff-choice, single-chooser rules that do not effectively include an element of defendant choice. Plaintiffs traditionally gain put options only after a defendant temporarily impairs their initial entitlement. Thus, an intentional taking by the defendant effectively converts such single-chooser put options into dual-chooser allocations. Only if the defendant takes unintentionally or if the plaintiff comes to a pre-existing nuisance might we have a circumstance that produces a true plaintiff-choice allocation.[81] But our motivation for studying dual-chooser allocations does not spring solely from their current common law existence — rather, we intend to show that dual-chooser rules at times are more allocatively efficient than single-chooser rules and should accordingly be added more consciously to policymakers' tool-kits.

A. *Selecting Optimal Damages*

Damages that maximize the allocative efficiency of plaintiff-presumption dual-chooser are those allocations such that:

$$f_\Pi(v_\Pi = D_{\Pi P})E[v_\Delta - D_{\Pi P} \mid v_\Delta > D_{\Pi P}] = f_\Delta(v_\Delta = D_{\Pi P})E[D_{\Pi P} - v_\Pi \mid D_{\Pi P} > v_\Pi] \qquad (C)$$

remembering that v_Π and v_Δ represent the plaintiff's and defendant's respective valuation, $f_\Pi(v_\Pi)$ and $f_\Delta(v_\Delta)$ represent the plaintiff's and defendant's respective probability distribution of these valuations, and E represents an expectations operator.[82] While this optimization equation (which economists call a "first-order condition") initially seems forbidding, it has a straightforward intuition. By setting damages, the court simultaneously determines two allocative margins: the margin on which the plaintiff will veto allocations of the entitlement to the de-

---

81. In contrast, it is usually possible to implement single chooser rules in which the defendant has the sole allocative choice. For example, the defendant's choice both to commence pollution and to exercise its call-option amount to a unitary decision to take the entitlement and pay damages. Of course, plaintiffs often make deliberate choices about whether to produce the entitlement in the first place — so the plaintiff's initial investment decision may introduce elements of plaintiff choice into what otherwise would be a defendant chooser rule.

There may, however, also be some limits on the ability of courts to vary how the payoff is distributed. If polluting potentially exposes the polluter to a pay-or-pay allocation, then a polluter who will ultimately choose not to allocate the entitlement to itself will prefer not to pollute in the first place. *See* Levmore, *supra* note 9, at 2168-70 (discussing impact of expected damages on litigants' initial incentive to pollute or bring suit).

82. *See* Appendix, *supra* note 40. Solving an analogous implicit formula for D yields the optimal damages for the defendant-presumption dual-chooser allocation:

$$f_\Pi(v_\Pi = D_{\Pi P})E[D_{\Delta P} - v_\Delta \mid D_{\Delta P} > v_\Delta] = f_\Delta(v_\Delta = D_{\Delta P})E[v_\Pi - D_{\Delta P} \mid v_\Pi > D_{\Delta P}].$$

fendant, and the margin on which the defendant will veto allocations of the entitlement to the defendant. The court's task is to find the single damage amount that optimally trades off these two veto margin effects. The damage amount that solves the foregoing equation accomplishes precisely this task.

A small increase in the size of $D_{\Pi P}$, from D' to D'', for example, has the simultaneous effect of (i) decreasing the plaintiff's willingness to veto allocations to the defendant, and (ii) increasing the defendant's willingness to veto allocations to the defendant.[83] The left-hand side of the optimization equation describes the marginal impact on allocative efficiency of decreases in the plaintiff's willingness to veto. It represents the changed likelihood that the plaintiff will veto, $f_\Pi(D_{\Pi P})$, multiplied by the consequent expected change in allocative efficiency, $E[v_A - D_{\Pi P} \mid v_A > D_{\Pi P}]$. The right-hand side of the equation analogously represents the marginal impact on allocative efficiency of increasing the defendant's willingness to veto defendant allocations. Solving the equation for $D_{\Pi P}$ therefore yields the damage amount that equates and thereby balances the marginal effects on allocative efficiency from the simultaneous increases in both the plaintiff's willingness to accept and the defendant's willingness to veto allocation of the entitlement to the defendant.

In practice, the optimal damage amount will often represent a weighted average of the plaintiff's and the defendant's mean valuations. This makes some intuitive sense when one remembers that setting damages equal to the litigants' mean valuations optimizes efficiency whenever a single litigant can veto and thereby block the allocation of an entitlement. It should therefore not be surprising that the optimal joint veto amount will normally lie somewhere between the optimal damage amounts of the two single-chooser allocations. Moreover, while it outstrips the quality of the evidence, as well as judicial temperament and training, to ask judges to solve explicitly the foregoing equation for an optimal D, the idea that optimal plaintiff-presumption dual-chooser damages will often be close to the average of the litigants' mean valuations provides some pragmatic guidance in setting damages.[84]

---

83. With higher damage amounts, the plaintiff becomes less likely to veto a defendant allocation because the plaintiff becomes more willing to allocate the entitlement to the defendant as long as:

$$v_\Pi < D_{\Pi P}.$$

Higher damage amounts increase the defendant's likelihood of vetoing defendant allocations because they makes it less likely that:

$$D_{\Pi P} < v_A.$$

84. *See* Ayres, *supra* note 9, at 827 (discussing possibility of a "joint-veto" regime with damages that "will approximately split the difference between the Resident's and the Polluter's mean valuations").

We can, however, conjure up examples where the optimal dual-chooser damages do not lie between the mean valuations of the litigants. For example, if the plaintiff's valuation is uniformly distributed between $25 and $75 and the defendant's is uniformly distributed between $0 and $100, the optimal plaintiff-presumption damages equal $62.50, an amount higher than the litigants' mean valuation of $50.[85] This happens because the plaintiff is a systematically poorer chooser than the defendant. Because the plaintiff's valuation in this example varies less than the defendant's valuation, the defendant has a greater informational advantage, relative to the court, than the plaintiff. This means that a defendant single-chooser allocation would be more efficient than a plaintiff single-chooser allocation in this case. The defendant's greater informational advantage also affects how an efficiency-minded court would distribute allocative power between the litigants in tailoring a dual-chooser rule. When the plaintiff has less private information and thus makes a relatively poor chooser, damage amounts that exceed both parties' mean valuations lower the odds that the plaintiff will veto a defendant allocation. Erring on the side of more frequent defendant vetoes yields greater allocative efficiency than more frequent plaintiff vetoes because the defendant makes more informed, and therefore more efficient, choices than does the plaintiff. We will return to this example again when we compare the four basic allocations. As we will see, the same effect that causes dual-chooser damages to take on extreme, and seemingly counterintuitive, values also causes single-chooser allocations to be more efficient than dual-chooser allocations.

B.   *Selecting the Distribution*

Before moving on to compare the relative allocative efficiency of the different rules, we pause here to comment on payoff distribution under the dual-chooser allocations. As with the single-chooser allocations, an infinity of dual-chooser implementations permits courts to vary the distribution of payoffs among the individual litigants while leaving allocative efficiency unaffected. Just as we had earlier distinguished between damages paid for self-allocation of an entitlement, D_C, and damages paid for allocation to the nonchooser, D_{NC}, under the plaintiff-presumption dual-chooser rules one can also distinguish between:

---

85. Applying our earlier equation (C) for optimal dual-chooser damages, *see supra* note 82 and accompanying text, to the particular uniform distributions in this example implies that optimal damages will solve the following equation:

$$\frac{1}{50}\int_{D}^{100}(v_\Delta - D)\frac{1}{100}dv_\Delta = \frac{1}{100}\int_{25}^{D}(D - v_\Pi)\frac{1}{50}dv_\Pi,$$

which, when solved, yields $D = \$62.50$.

$D_{No\ Veto}$ =   the price the defendant pays the plaintiff if neither litigant vetoes the allocation of the entitlement to the defendant; and,

D_{Veto} =   the price the defendant pays the plaintiff if one of the litigants vetoes a defendant allocation so that the entitlement goes to the plaintiff.

If we label as D_{IIP} the damage amount that solves the foregoing optimization equation (C), then a continuum of allocatively identical plaintiff-presumption rules emerges that take the form:

$$D_{No\ Veto} - D_{Veto} = D_{IIP}.$$

For example, if the allocatively optimal damages equal $50, then a dual-chooser allocation that required the defendant to pay $25 to the plaintiff when the entitlement is allocated to the defendant ($D_{No\ Veto}$ = $25), and required the plaintiff to pay $25 to the defendant when the entitlement is allocated to the plaintiff (D_{Veto} = −$25), would produce exactly the same allocative choices by the litigants as a simpler rule that required a payment of $50 by the defendant if the parties jointly chose to allocate the entitlement to the defendant. In either case, the plaintiff would choose to allocate the entitlement to the defendant only if the plaintiff valued it less than $50, and the defendant would choose to allocate the entitlement to itself only if the defendant valued it more than $50.

Thus, as before, a lock-step increase or decrease in $D_{No\ Veto}$ and D_{Veto} does not affect the litigants' allocative choice and a fortiori the allocative efficiency. Rather, such changes have the effect of increasing or decreasing the plaintiff's average share of the expected total payoff. Courts can therefore freely distribute the total expected payoff between the litigants as they see fit. Although computationally more difficult, a dual chooser analog to the "alpha" implementations of the single-chooser allocations would also allow courts to allot a fraction α of the expected joint payoffs to the plaintiff and a fraction $(1 - \alpha)$ to the defendant.[86]

---

86. For example, assume that D_{IIP} equals the optimal damages for the plaintiff-presumption dual-chooser allocation. Assume further that under the straightforward "double-put" regime (discussed *supra* notes 75-76 and accompanying text, in which $D_{No\ Veto} = D_{IIP}$ and $D_{Veto} = 0$), the expected payoffs for the plaintiff and the defendant equal EP_d and EP_{II}, respectively. One can then show that an implementation that sets $D_{No\ Veto} = D_{IIP} - EP_{II}$ and $D_{Veto} = - EP_{II}$ would distribute the entire joint expected payoff to the defendant; and an implementation that sets $D_{No\ Veto} = D_{IIP} + EP_d$ and $D_{Veto} = EP_d$ would distribute the entire joint expected payoff to the plaintiff. More generally, an implementation that sets $D_{No\ Veto} = D_{IIP} - \alpha EP_{II} + (1 - \alpha)\ EP_d$ and $D_{Veto} = -\alpha EP_{II} + (1 - \alpha)EP_d$ would distribute a fraction α of the joint expected payoff to the plaintiff and a fraction $(1 - \alpha)$ to the defendant.

C.   *Selecting the Choosers*

As an initial task in identifying the more efficient chooser, a court must determine which of the two dual-chooser allocations is more efficient, the plaintiff-presumption allocation or the defendant-presumption allocation. Graphically, this comes down to determining whether to apportion the quadrants II and III (of Figure 4 below) to the plaintiff or to the defendant. Both dual-chooser allocations apportion the first quadrant to the defendant and the fourth quadrant to the plaintiff. But the plaintiff-presumption regime apportions both of the remaining quadrants, II and III, to the plaintiff, while the defendant-presumption regime apportions the off-diagonal quadrants to the defendant.[87]

Figure 4
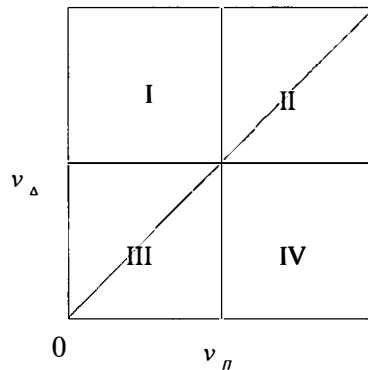
It turns out that the plaintiff-presumption allocation tends to be more efficient than the defendant-presumption equilibrium when:

$$\mu_\Pi > \mu_\Delta.^{88}$$

For example, imagine that the defendant's valuation is uniformly distributed between $0 and $100, and that the plaintiff's valuation is uniformly distributed between $80 and $180. The optimal damages for either a plaintiff-presumption or a defendant-presumption allocation will equal $90, the average of the litigants' mean valuations.[89] But, as

---

87. Besides this "off-diagonal" effect, the two types of dual-chooser rules can produce different optimal damages, which might affect which equilibrium generates more allocative efficiency.

88. For example, when the litigants' valuations are uniformly distributed, a higher plaintiff mean implies that the plaintiff-presumption equilibrium will produce a higher expected joint payoff than the defendant-presumption equilibrium.

89. Applying our earlier equation (C) for optimal dual-chooser damages, *see supra* text accompanying note 82, to the particular uniform distributions in this example implies that optimal damages will solve the following equation:

Figure 5 shows, a plaintiff-presumption allocation is much more effi-
cient than a defendant-presumption allocation because it apportions
the off-diagonal quadrants (II and III) to the plaintiff. Figure 5 shows
that much more of the probability mass in quadrants II and III lies be-
low the first-best forty-five degree line than above it.[90] In other words,
the plaintiff's valuation will probably exceed the defendant's valuation
in these quadrants. This simple example provides intuitive support for
the claim that the plaintiff-presumption allocation will tend to domi-
nate the defendant-presumption equilibrium when the plaintiff's mean
is higher than the defendant's. The converse also tends to hold true.

Figure 5

Entitlement allocated to defendant

Entitlement allocated to plaintiff
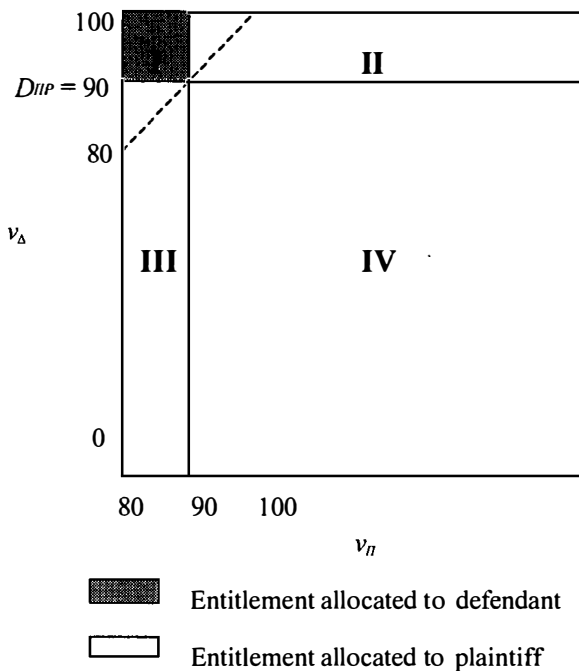
$$\frac{1}{100}\int_D^{100}(v_\Delta - D)\frac{1}{100}dv_\Delta = \frac{1}{100}\int_{80}^D (D - v_\Pi)\frac{1}{50}dv_\Pi,$$

which when solved yields D = $90.00.

90. The forty-five degree diagonal line represents the first-best allocation, possible only
with perfect information.

Appreciating the theoretical and practical possibility of rules partially rehabilitates the Restatement's focus on the parties' relative valuations. Although the relative size of the litigants' mean valuations does not significantly bear on the selection between the two single-chooser allocations, it does bear not only on the selection between dual-chooser allocations, but also (as we are about to see) on the selection between dual- and single-chooser allocations themselves.

D. *Selecting Among the Four Foundational Single-Price Allocations*

To summarize our results thus far, we have shown that four basic allocations exist — two single-chooser and two dual-chooser allocations. We have derived the optimal damages for each and shown that courts can divide the expected total payoff between the litigants as it wishes without affecting the allocative efficiency of any of the four allocations. Finally, we have shown that under a single-chooser rule, courts should appoint as chooser the litigant with the more speculative valuation. Under a dual-chooser rule, on the other hand, courts should allow either party to veto allocations to the party with the lower mean.

We can restate these last two results more mathematically as follows:

(1')   when $\sigma_\Pi > \sigma_\Delta$, $E[TP]_\Pi > E[TP]_\Delta$, and

(2')   when $\mu_\Pi > \mu_\Delta$, $E[TP]_{\Pi P} > E[TP]_{\Delta P}$,

where $E[TP]$ represents the expected total payoff, and the subscripts Π, Δ, ΠP, and ΔP refer to the plaintiff-choice, defendant-choice, plaintiff-presumption dual-chooser and defendant-presumption dual-chooser allocations, respectively. In the following discussion, we refer to the first result as the "different variance" effect and to the second result as the "different mean" effect.

Our final task is to try to identify which of the four allocations is likely to maximize allocative efficiency. In other words, we need to assess when the better single-chooser allocation dominates (or is dominated by) the better dual-chooser allocation. As an initial matter, there are many circumstances in which the four basic equilibria will be allocatively equivalent. For example, if the litigants have identical and symmetric probability distributions of valuation, the four classes of rules will generate identical damages and identical expected total payoffs. Under such circumstances, apportionment of quadrants II and III to either litigant will be equally efficient, and so the four different permutations of quadrant allocations generated by the four basic allocations will be equivalent.

But it turns out for more general probability distributions that the dual-chooser allocations are sometimes better than, and at other times inferior to, single-chooser allocations. Two heads can be better than

one; but sometimes one head — that is, a single allocative chooser — can be better than two.

The optimal allocation will turn on the relative size of the different-mean and different-variance effects discussed above. When the difference in valuation means is large relative to the difference in valuation variances, a dual-chooser rule that allows either party to veto an allocation to the lower valuing litigant will tend to dominate the other three basic allocations. In contrast, when the difference in valuation variances is large relative to the difference in valuation means, a single-chooser rule that allows the higher variance litigant to allocate the entitlement will tend to dominate the other three basic allocations.[91]

Put more intuitively, two choosers will tend to be better than one when the litigants have relatively different means but relatively similar variances. A sufficiently wide divergence in the litigants' means undermines the efficiency of a single-chooser rule because no single-chooser has sufficient information to make a nuanced allocation. By "nuanced allocation," we mean nothing more than an equilibrium in which the chooser makes more refined decisions about who will be allocated the entitlement — at times allocating the entitlement to each litigant. As the different-mean effect comes to dominate the different-variance effect, however, single choosers will tend to make "less-nuanced" decisions — allocating the entitlement to the same litigant all the time. Under such conditions, a dual-chooser rule can increase allocative efficiency, because the litigants' combined allocative power can identify when a given allocation should occur even when single choosers would not have chosen it.

For example, consider again a defendant whose value is uniformly distributed between $0 and $100 and a plaintiff whose value is uniformly distributed between $80 and $180, as in Figure 5. Notice that we have intentionally constructed the example so that the different-mean effect will dominate the different-variance effect. The variances of the litigants' values match ($\sigma^2_\Pi = \sigma^2_\Delta = 833.33), but their means differ ($\mu_\Pi = $130 > \mu_\Delta = 50). Our prior analysis indicated that the equal variances would tend to make the plaintiff-choice rule allocatively equivalent to the defendant-choice rule. Moreover, the unequal mean valuations tend to favor the plaintiff-presumption dual-chooser alloca-

---

91. For example, for uncorrelated, uniform distributions it can be shown that when there is no mean difference but only a variance difference, the better single-chooser rule dominates the better dual-chooser rule and hence the better single-chooser rule is globally optimal (within the class of single-price liability rules). Conversely, when there is no variance difference, but only a mean difference, it can be shown that the better dual-chooser rule dominates the better single-chooser rule and hence the better dual-chooser rule is globally optimal.

tion over the defendant-presumption dual-chooser allocation.[92] But what still needs to be determined is whether the plaintiff-presumption dual-chooser allocation dominates either of the equivalent single-chooser allocations.

As suggested above, it does, and here's why. Either of the single-chooser allocations produces a non-nuanced decision — allocating the entitlement exclusively to the plaintiff. If the plaintiff acts as chooser and damages equal the defendant's mean of $50, then the plaintiff will always allocate the entitlement to itself.[93] Conversely, if the defendant acts as chooser and damages equal the plaintiff's mean valuation of $130, the defendant will never allocate the entitlement to itself.[94] When the difference in mean valuations becomes sufficiently large, neither chooser possesses information sufficient to make a nuanced allocation.

A dual-chooser allocation, by contrast, improves allocative efficiency because the joint choice of the litigants allows the entitlement to go at times to the defendant when the defendant values the entitlement more. The dual-chooser allocation allows the low mean litigant (here, the defendant) to end up with the entitlement when it has an unusually high value and when the high mean litigant (the plaintiff) has an unusually low value. As discussed above,[95] the optimal dual-chooser damages for this example equal $90. With a single-chooser allocation, a damage amount of $90 would not produce an efficient outcome on average because such an amount would cause the lone chooser to allocate the entitlement inefficiently to the defendant too often. The plaintiff-presumption allocation, however, can efficiently operate with a damage level of $90 because the entitlement will go to the defendant only when the defendant's valuation exceeds $90 and the plaintiff's valuation falls short of $90. Figure 5 depicts this more nuanced allocation. When both litigants have a say in the ultimate allocation, the expected total payoff increases.[96]

---

92. The example can also be easily converted to show how a defendant-presumption rule can dominate any of the other three foundational allocations by simply reversing the identity of the plaintiff and defendant probability distributions — so that the plaintiff's value would be distributed uniformly between $0 and $100, and the defendant's value between $80 and $180.

93. This happens because the plaintiff's valuation ranges from $80 to $180 and thus will always exceed the $50 damage payment for allocating the entitlement to the defendant.

94. Likewise, this happens because the defendant's valuation ranges from $0 to $100 and thus the $130 damage payment for allocation to the plaintiff would always exceed the defendant's valuation.

95. *See supra* note 89 and accompanying text.

96. Expected joint profits under a plaintiff-presumption allocation equal $130.10, as compared to single-chooser joint profits of $130.00. When the single-chooser rules allocate the entitlement exclusively to one litigant, the dual-chooser rules can never produce a lower expected total payoff because we can set dual-chooser damages to produce an always non-

When the litigants' estimated mean valuations differ greatly, the superiority of the dual-chooser allocations resolves an embarrassing disconnect between our liability rule analysis above and cases like *Boomer*. In *Boomer*, the court estimated that the Polluter valued the right to pollute much more than the Residents valued the right to no pollution.[97] Simple single-chooser rules of either the call or put varieties (Rules 2, 4, 5 and 6) are allocatively embarrassing because setting the damages equal to the nonchooser's mean would produce non-nuanced equilibria that are allocatively indistinguishable from simply giving the polluter a property-rule interest in the entitlement (Rule 3). *Boomer*-like residents confronting an allocative price set equal to the polluter's high mean value would never allocate the entitlement to themselves, and an *Atlantic Cement*-like polluter confronting an allocative price set equal to the residents' low mean value would never fail to allocate the entitlement to itself. When the litigants' means vary greatly, single-chooser rules are allocatively a non-event.[98] In contrast, the dual-chooser rules intentionally seek out a damage amount between the litigants' mean valuations — where the litigants' probability distributions overlap — so as to produce more efficient, nuanced allocations. Unlike the single-chooser allocations, dual-chooser allocations do not force the court to set damages only at the plaintiff's or the defendant's mean valuation.[99]

Dual-chooser allocations, however, are not always preferable. When the litigants have relatively divergent variances but similar

---

nuanced equilibrium. For example, if the plaintiff-presumption damages were set at $130, then the defendant would always veto allocating the entitlement to itself.

97. Boomer v. Atlantic Cement Co., 257 N.E.2d 870, 875 (N.Y. 1970). *Whalen v. Union Bag & Paper Co.*, 101 N.E. 805 (N.Y. 1913), also provides an extreme example of the willingness of some courts to issue injunctions where a defendant's loss far exceeds the plaintiff's benefit. In the *Union Bag* case (discussed later in *Boomer*, 257 N.E.2d at 872) the court assessed the plaintiff's harm at $100 per year, while compliance with the injunction caused the permanent closing of a mill at an investment loss of more than $1,000,000. 101 N.E. at 805. The New York Court of Appeals concluded, "[A]lthough the damage to the plaintiff may be slight as compared with the defendant's expense of abating the condition, that is not a good reason for refusing an injunction." *Id.* at 806; *see also* Ian Ayres & Kristin Madison, *Threatening Inefficient Performance of Injunctions and Contracts*, 148 U. PA. L. REV. 45 (1999).

98. The single-chooser rules can also be distributively embarrassing when there is a great disparity in the litigants' mean valuations. Such rules can force the policy makers at times to confront unpalatable divisions. For example, if a court decided that the residents in *Boomer* were the more efficient choosers, it would have to choose between a call rule, which might force residents to pay a large sum to stop the pollution, and a put rule, which might force the polluter to disgorge all of its expected profits from polluting. Our foregoing convexity result, however, resolves the distributive embarrassment under either the single- or dual-chooser implementations.

99. Of course, if their mean valuations diverge too much (or more specifically, if their probability distributions do not overlap), then the court cannot achieve, nor would they desire, a more nuanced allocation, even under the dual-chooser rules. In such an extreme case, the court would know for certain that one litigant's valuation exceeded the other litigant's valuation and would thus face no informational disadvantage. The court could therefore simply award the entitlement to the higher valuing litigant.

mean valuations, greater allocative efficiency would be produced by a single-chooser allocation — granting the allocative choice to the chooser with the larger variance. The intuition here is that one head can be better than two when one of the heads is relatively ignorant. When the litigants' valuation variances differ greatly, courts should exclude the low variance litigant from having any influence on the ultimate allocation. A litigant with a relatively small variance has very little informational advantage over either the court or the opposing litigant. When the different-variance effect thus dominates the different-mean effect, courts should select a single chooser rather than allow a second litigant with very little private information to muck up the allocation.

In a recent article on the "anticommons," Michael Heller has persuasively argued that there is an efficient number of vetoers to the deployment of entitlements.[100] The traditional commons problem occurs is when the number of potential vetoers is too small — no one can veto the deployment of the entitlement. But an analogous anticommons problem exists when there are too many vetoers who can block the deployment of an entitlement. Heller's canonical examples concern modern Russia, where entrepreneurs may need the approval of dozens of organizations before they can use a piece of property to open a business.[101] Our finding that single-chooser allocations can dominate dual-chooser allocations suggests that the anticommons inefficiency can kick in even when the number of vetoers is just two. Sometimes it is efficient to have only a single vetoer.

For a more concrete example, consider again a defendant whose value is uniformly distributed between $0 and $100 and a plaintiff whose value is uniformly distributed between $25 and $75. Notice now that this example in intentionally constructed so that the different-variance effect will dominate the different-mean effect, for the litigants have the same mean valuations ($\mu_\Pi = \mu_\Delta = \$50$), but their valuation variances differ ($\sigma^2_\Delta = \$833.33 > \sigma^2_\Pi = \$208.33$). As our prior analysis showed, the equality of means will tend to make the two dual-chooser rules allocatively equivalent. The inequality in variances also suggests that the defendant-choice allocation will dominate the plaintiff-choice allocation.[102] But what still needs to be shown is why the defendant-choice (single-chooser) allocation dominates either of the equivalent dual-chooser allocations.

---

100. Michael A. Heller, *The Tragedy of the Anticommons: Property in the Transition from Marx to Markets*, 111 HARV. L. REV. 621 (1998).

101. *See id.* at 633-42.

102. The example can also be easily converted to show how a plaintiff-joint-veto rule can dominate any of the other three foundational allocations by simply reversing the identity of the plaintiff and defendant probability distributions — so that the plaintiff's value was distributed uniformly between $0 and $100, and the defendant's value between $80 and $180.

When litigant variances differ greatly, dual-chooser allocations become problematic because they allow the more informationally disadvantaged chooser — the chooser with the less variant valuation — to have too much control over the entitlement's ultimate allocation. For example, as shown in Figure 6's depiction of the plaintiff-presumption allocation, even when the plaintiff's value is higher than the nominal damage price, this does not provide very good assurance that the plaintiff's value is higher than the defendant's (higher variance) valuation.

Indeed, as discussed above, the relative inferiority of the plaintiff as chooser in this example causes the optimal damages for the plaintiff-presumption rule to exceed the mean valuations of both parties, such that $D_{\Pi P} = \$62.50 > \mu_{\Pi} = \mu_A = \$50$. Because the plaintiff-presumption rule gives the plaintiff too much power to self-allocate the entitlement, the optimal damages under this rule are raised to reduce the plaintiff's incentives to self-allocate. Because the litigants also differ in their level of private information, it is better to induce the defendant to inefficiently under-take (that is, to allocate the entitlement too often to the plaintiff) in order to restrict the poorly informed plaintiff's incentive to over-take (that is, to self-allocate the entitlement too often). The fact that optimal dual-chooser damages lie outside the litigants' mean valuations is evidence that one of the litigants is a particularly inefficient chooser and therefore provides evidence that dual-chooser allocations will tend to be inefficient relative to single-chooser allocations.

The graphical difference between these allocations concerns which litigant will ultimately possess the entitlement in quadrant II. The defendant-choice single-chooser rule allocates this quadrant to the defendant, thereby enhancing allocative efficiency. Witness Figure 6, which shows that more of the probability space in quadrant II contains situations in which the defendant's valuation exceeds the plaintiff's (as seen from the larger mass of this quadrant above the forty-five-degree first-best efficiency line). The simpler defendant-choice rule divests the plaintiff of allocative power and thus lowers the optimal damages back down to $50 and increases the expected joint payoff from $60.54 under the plaintiff-presumption dual-chooser allocation to $62.50 under the defendant-choice allocation.[103]

---

103. An analogous argument also shows why, in this example, a defendant-choice allocation dominates the defendant-presumption dual-chooser regime. Under a defendant-presumption rule, the defendant cannot prevent the plaintiff from *under*-taking and putting the entitlement to the defendant even when the high expected mean defendant has a seriously low actual valuation. This scenario exists in much of the lower-left quadrant. A defendant-presumption rule lowers the optimal damages below the litigants' mean valuations in order to increase the poorly informed plaintiff's incentive to allocate the entitlement to the defendant unilaterally. Rather than distort damages this way, however, we can even more efficiently give all the allocative power to the defendant through a single-chooser rule.

Figure 6

Entitlement allocated to defendant

Entitlement allocated to plaintiff
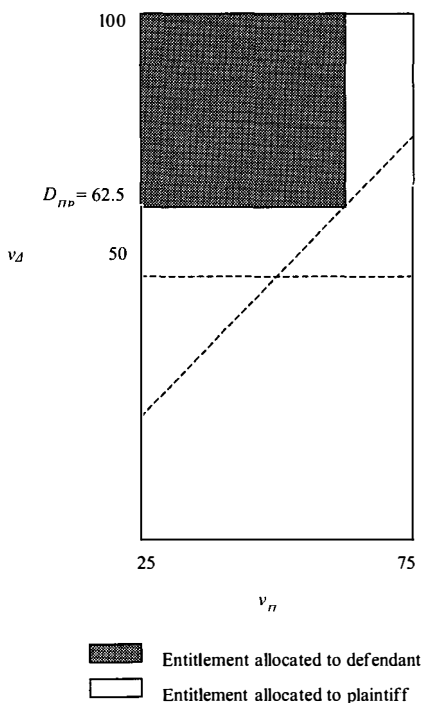
The foregoing two examples were intentionally constructed to highlight the impact of the different-mean and different-variance effects by alternately setting one or the other of these differences to zero. But, more generally, the litigants' distributions may display a mixture of both mean and variance differences. In such cases, selection of the most efficient allocation will turn on which of these difference effects dominates. Figure 7 graphically shows for a uniform example which of the four basic allocations will dominate — holding variance constant — as we change the relative mean valuations of the two litigants. For example, in panel (a), imagine that the plaintiff's valuation is uniformly distributed over a $100 range, while the defendant's valuation is uniformly distributed over only a $50 range. Of course, this assumption implies that the variance of the plaintiff's value exceeds the variance of the defendant's value. We can then add the litigant's relative mean valuations to the figure by drawing in a particular forty-five-degree locus of points in which the litigants' valuations are equal. Thus, the forty-five-degree line toward the upper left-hand corner depicts an example in which the plaintiff's mean is greater than the defendant's, while the forty-five-degree line toward

the lower-right hand corner depicts an example in which the defendant's mean is greater than the plaintiff's.

Figure 7

(a) Example where $\sigma_\Pi > \sigma_\Delta$ (b) Example where $\sigma_\Delta > \sigma_\Pi$
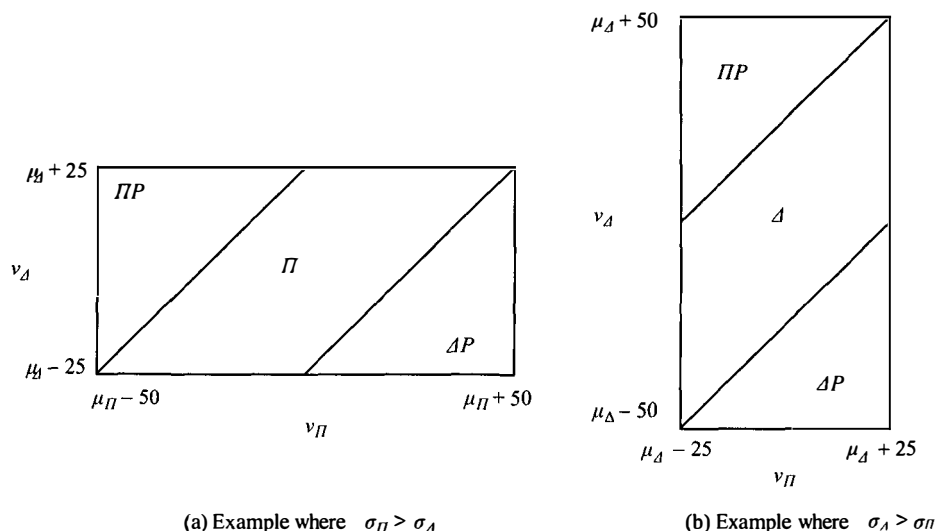
Panel (a) shows that — holding the variance difference constant — as the mean difference varies (from a higher plaintiff mean to a higher defendant mean), three different fundamental allocations in turn become the most efficient. For sufficiently large differences in litigant means in the plaintiff's favor, the different-mean effect dominates the different-variance effect, making the plaintiff-presumption dual-chooser allocation the most efficient.[104] For intermediate differences in the litigants' means, the different-variance effect dominates, and so the plaintiff-choice single-chooser allocation dominates. Finally, for sufficiently large mean differences in the defendant's favor, the different-mean effect again dominates the variance effect, this time rendering the defendant-presumption dual-chooser allocation the most efficient. Panel (b) works out the optimal allocations for an example in which the variance of the defendant's valuation exceeds that of the plaintiff's. When the different-mean effect dominates, one of the dual-chooser allocations is most efficient; when the difference in the litigants' means is not as large, the variance effect dominates, and a

---

104. If $\mu_\Pi - \mu_\Delta > \$25$, then given the \$50 assumed difference in the litigants' uniform distribution limits, the plaintiff-presumption allocation produces the highest payoffs. If $\$25 > \mu_\Pi - \mu_\Delta > -\$25$, then the plaintiff-choice single-chooser allocation produces the highest payoffs. Finally, if $\mu_\Pi - \mu_\Delta < -\$25$, then the defendant-presumption allocation produces the highest payoffs.

defendant-choice single-chooser allocation is most efficient. Indeed, for the uncorrelated, uniform distribution, the best single-chooser allocation will at least weakly dominate the best dual-chooser allocation whenever:

$$\left| \lambda_\Pi - \lambda_\Delta \right| \geq \left| \mu_\Pi - \mu_\Delta \right|$$

(where the plaintiff and defendant valuations are uniformly distributed between $\mu \pm \lambda$).[105]

In sum, we have now provided a complete analysis of the four basic single-price allocations. We have also shown what damages optimize allocative efficiency for each type of allocation. We have shown when each type of the four allocations will tend to be allocatively optimal. And we have shown how courts are free to independently distribute the expected payoff between the litigants.

III. SECOND (AND HIGHER) ORDER RULES

In a previous article, Ian Ayres and Jack Balkin analyzed a type of regime under what they called a "second-order" liability rule.[106] A second-order rule constitutes a type of dual-chooser allocation because each litigant makes successive allocative choices. Unlike the dual-chooser allocations analyzed above, however, the second-order rules do not rely on a single judicially determined allocative price. Instead, second-order rules force one litigant to confront one allocative exercise price, and force the other litigant to confront a different allocative exercise price.

For example, if both the plaintiff and defendant have valuations uniformly distributed between $0 and $100, one type of optimal second-order liability rule would induce the defendant to exercise an option to take when its value is higher than $33.33, but would induce the plaintiff to take back the entitlement when its entitlement is greater than $66.66. Ayres and Balkin showed that second-order rules mimicked the allocations of an auction with minimum bid increments.[107] This allocative equilibrium is graphically displayed in Figure
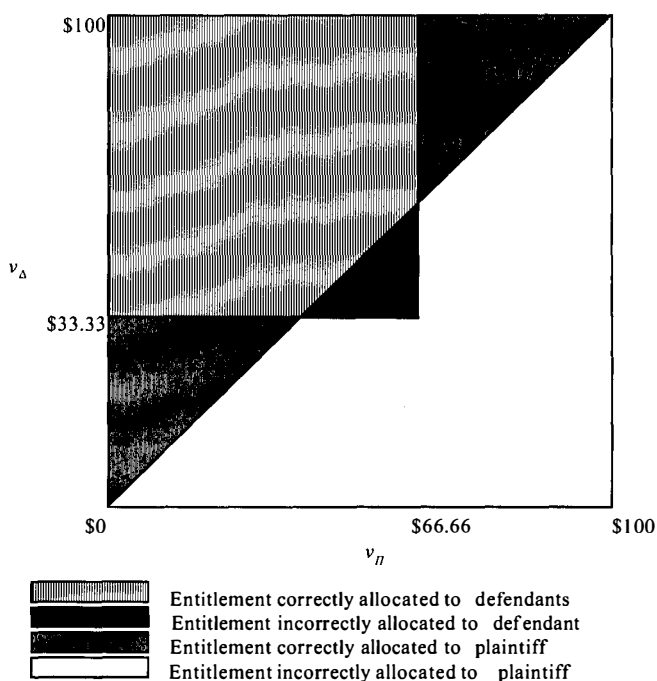
---

105. *See* Avraham, *supra* note 22; Appendix, *supra* note 40.

106. *See* Ayres & Balkin, *supra* note 9.

107. *See id.* at 727-28. Unlike traditional auctions, however, in which the proceeds of the auction go to a non-bidder, the *internal* auction produced by second (and higher-) order rules is one in which the proceeds from the winning bidder are paid to the losing bidder. One could imagine more explicit (external) auction implementations — in which the state initially took an entitlement (for some fixed payment to plaintiff) and then auctioned the entitlement to the highest bidder (with the state retaining auction proceeds). Such an auction would fully economize on the litigants' private information and could be implemented with much more transparent bidding strategies. This more explicit auction might prove the most efficient of all depending on the relative transaction costs. And by varying the initial payments from the government to each of the litigants, the explicit auction approach is also compatible with the type of distribution decoupling discussed above.

8. The plaintiff-presumption dual-chooser allocation is a type of second-order liability rule, but one in which the first- and second-order exercise prices are *constrained to be the same.*

Figure 8

Entitlement correctly allocated to defendants
Entitlement incorrectly allocated to defendant
Entitlement correctly allocated to plaintiff
Entitlement incorrectly allocated to   plaintiff

Ayres and Balkin also showed that allowing exercise prices to vary induces strategic taking behavior by the initial decisionmaker. For example, in the foregoing hypothetical a defendant with a private valuation of only $33 would still find it advantageous to exercise an initial call option to take an entitlement for $33.33 (thus, paying more than its valuation) because doing so might induce the plaintiff to take back and pay the defendant $66.66. In order to induce optimal taking strategies, the nominal damage amounts would need to be set at $44.44 (for the defendant's call option) and $66.66 (for the plaintiff's take-back call option). Given the odds that the plaintiff's valuation will exceed $66.66 and that the plaintiff will therefore take the entitlement back, the defendant's expected profit from strategically taking the entitlement for a price of $44.44 is negative when the defendant's true valuation falls anywhere below $33.33. Because of this incentive for strategic taking, the exercise prices of second-order rules do not

induce transparently obvious player strategies.

Like its constrained dual-chooser siblings, second-order rules have multiple option variations that will induce allocatively equivalent allocations. For example, either of the following allocations produces what we call an "unconstrained-plaintiff-presumption" allocation:

$$\Pi: \quad E - \text{call(Low)} + \text{call(High)} \qquad \Delta: \quad + \text{call(Low)} - \text{call(High)} \qquad (1)$$

$$\Pi: \quad E + \text{put(High)} - \text{put(Low)} \qquad \Delta: \quad - \text{put(High)} + \text{put(Low)}, \qquad (2)$$

where the parenthetical *Low*'s and *High*'s merely indicate the ordinal size of the two exercise prices.

In the first implementation (1), the plaintiff (a) receives the initial entitlement, (b) subject to the defendant's call option, but (c) also has a call option to take back the entitlement at a higher price. Like the earlier plaintiff-presumption dual-chooser allocation, this allocation allows either litigant to veto the transfer of the entitlement to the defendant. Defendants can veto the transfer of the entitlement to themselves by refusing to exercise their initial call option, and even if the defendants exercise their call options, the plaintiffs can veto the transfer by exercising their take-back call option. Alternatively, under the second implementation (2), the plaintiff receives (a) the initial entitlement and (b) an option to put to the defendant at a high exercise price, but (c) the defendant receives an option to put the entitlement back to the plaintiff at a lower exercise price.[108] An analogous second-order allocation can be implemented by what we call "unconstrained defendant-presumption" allocations simply by reversing the foregoing plaintiff and defendant roles.

While these second-order rules may seem the most bizarre of our ever expanding universe of "liability" rules, Ayres and Balkin have shown that the court implicitly implemented a second-order rule in the venerable common law chestnut, *Vincent v. Lake Erie Transportation Co.*[109] In that case, the Minnesota Supreme Court held a ship owner liable when his ship damaged a dock while he attempted to moor the ship during a storm. The court simultaneously acknowledged that it would have ordered the dock owner to pay damages to the ship owner if the dock owner had subsequently unmoored the defendant's ship, causing it to suffer damage in the storm. As Ayres and Balkin noted:

---

108. The same allocation might also be induced by:

$$\Pi: - \text{put(Low)} + \text{call(High)} \qquad \Delta: E + \text{put(Low)} - \text{call(High)} \qquad (3)$$

Under this third implementation (3), the defendant has the initial entitlement and a put option with a low exercise price, while the plaintiff receives a call option with a high exercise price. The plaintiff and defendant might have to exercise their respective options simultaneously. The entitlement would end up with the defendant only if neither litigant vetoed such an allocation by exercising an option. Although the litigants might simultaneously exercise these options, we would still expect nontransparent strategies, as each party would hedge on its decision to exercise and hope that the other litigant would exercise at a more favorable price.

109. 124 N.W. 221 (Minn. 1910).

> *Vincent* [is] a vivid example of how the common law protects an option to take an entitlement (a liability rule) with another liability rule. The dock owner holds the initial entitlement to the physical security of the dock. The shipowner (because of the exigencies of the storm) has a first-stage option to "take" the dock by mooring the ship to it and by paying damages for any injury that results. The dock owner has a second-stage option to unmoor the ship, but at a cost: The dock owner gives up a cause of action against the shipowner for damages and exposes himself to tort liability for any damages to the ship and its crew.[110]

Under *Vincent*, both the ship and dock owners had a say in the ultimate allocation of the dock, and while nothing indicated that the court intended to set allocatively optimal rules, it certainly did not constrain the potential damages in the first and second stages to the same amount (in contrast with the aforementioned single-price, dual-chooser allocations).

This Section extends Ayres and Balkin's previous analysis by showing more generally how courts should tailor second-order damages, as well as by showing that we can apply our convexity result to second-order rules and again decouple a court's distributive and allocative concerns. Finally, we compare these unconstrained dual-chooser allocations to the four single-price allocations. Intuitively, it should not be surprising that a less constrained rule should operate more efficiently than more constrained rules, and indeed, we find that, absent transaction costs, second-order rules dominate.

A. *Selecting Optimal Damages*

Calculating optimal damages is computationally difficult — even more so than solving the dual-chooser optimization equation outlined above. We readily admit that such calculations will outstrip the computational capacity and temperament of judges. Still, it is possible to make some progress on the underlying process that would guide a judge inclined to make such an allocation (possibly with the aid of a fairly simple piece of computer software).[111] First, it is useful to break the task into two parts: (1) identifying the optimal "pivot values" that

---

110. Ayres & Balkin, *supra* note 9, at 716. Jon Hanson and Matt Stowe first noticed this unusual aspect of *Vincent. See* Jon Hanson & Matt Stowe, Lecture Notes, Torts, Harvard Law School (Fall 1996) (on file with the *Yale Law Journal*). Robert Ellickson long ago also proposed a modification of nuisance rules that would amount to a common law implementation of a second-order rule. *See* Robert C. Ellickson, *Alternatives to Zoning: Covenants, Nuisance Rules, and Fines as Land Use Controls*, 40 U. CHI. L. REV. 681 (1973).

111. The real difficulty facing a judge involves reduction of a litigant's valuation to a probability density function. Once a judge accomplishes this (and interactive software already exists to help quantiphobes undertake just this task through a series of yes/no questions), a piece of software can complete the fairly trivial next step of numerically calculating the optimal damages. We guess that more complicated computational problems are surmounted every time someone dials a phone number.

define when a chooser will optimally allocate the entitlement to itself or allocate it to the other litigant; and (2) identifying the nominal damages that will induce the litigants to adopt particular pivot values as their equilibrium strategies. For example, in the hypothetical situation above, we set nominal damages at $44.44 and $66.66 so as to induce pivot values of $33.33 and $66.66.

The first task of identifying the optimal pivot value has a surprisingly simple intuition, which we call the "dispositive-takings principle." To see how this principle might work, consider again a second-order regime in which the defendant has a first-stage call option to take and the plaintiff has a second-stage option to take back at a higher price. While it makes superficial sense for a court to want to induce takings when the taker believes that its value is greater than the expected value of the takee, this approach ignores the fact that, in a second-order regime, many of the first-stage takings will not be dispositive — that is, they will not allocate the entitlement. Some of the first-stage takees will have values higher than the expected values of the taker, and they will protect their interests by taking back. The optimal pivot values should focus on the allocations of particular choosers that are expected to be final, or dispositive. Hence, optimal exercise prices for second-order rules should induce the first-stage taker to take whenever the taker's private valuation exceeds the expected valuation of takees *who will not take back.*[112] The dispositive-takings principle requires that dispositive takings must on average increase welfare.

While the dispositive takings principle has intuitive appeal, implementing it would require a judge to simultaneously solve a system of two equations for the two unknown pivot values:

$$Pivot_\Delta = E\{v_\Pi \mid v_\Pi < Pivot_\Pi\}$$
$$Pivot_\Pi = E\{v_\Delta \mid v_\Delta > Pivot_\Delta\}$$

where $Pivot_\Delta$ and $Pivot_\Pi$ represent the pivot values at which defendants and plaintiffs choose to exercise their taking and take-back options, respectively.[113] The first pivot equation asks what is the plain-

---

112. Ayres & Balkin, *supra* note 9, at 717, proved that the dispositive-taking principal maximized allocative efficiency when the litigants' valuations were uniformly distributed and conjectured that the principal would also hold true for more general distributions. In the Appendix, *supra* note 40, we prove this to be true.

113. The equations can be restated in terms of the explicit integrals and probability density functions as well:

$$Pivot_\Delta = \int_{Pivot_\Pi}^{\infty} v_\Pi \, f_\Pi(v_\Pi) dv_\Pi$$

$$Pivot_\Pi = \int_{0}^{Pivot_\Delta} v_\Delta \, f_\Delta(v_\Delta) dv_\Delta$$

tiff's expected value given that the plaintiff will not take back. It is this class of plaintiffs for whom the defendant's taking will be dispositive. The second pivot equation asks what is the defendant's expected value given that it initially exercised its option to take. It is this class of defendant for whom the plaintiff's taking back will be dispositive.

But even after undertaking the arduous task of calculating the optimal pivot values, a court would then — because of the first-stage strategic exercise incentive — still need to solve a second system of equations to determine what nominal damage amounts will induce the desired pivot strategies. Because second-stage takings are nonstrategic, the second-stage damages are identical to the second-stage pivot value:

$$D_\Pi = Pivot_\Pi.$$

But the first-stage damages (because of the problem of strategic first-stage taking) are more complicated:

$$D_\Delta = P(v_\Pi < Pivot_\Pi)Pivot_\Delta + P(v_\Pi > Pivot_\Pi)Pivot_\Pi.$$

Notice that first-stage damages equal a weighted average of the two optimal pivot values (where the weights concern the probability that the plaintiff's value falls above or below the plaintiff's take-back pivot value). For example, if the plaintiff and defendant have identical probability density functions that take on a quadratic form with zero probability mass at $0 and $100,[114] then it can be shown in applying the foregoing dispositive takings equations that the players' joint payoffs are maximized if the defendant takes when its valuation is greater than $37.80 and the plaintiff takes back when its valuation is greater than $62.20. And to induce these pivot values, it is necessary for a court to set first-stage damages for defendant taking equal to $45.62 and second-stage damages for plaintiff taking back equal to $62.20.[115] Thus, while optimal damages for second-order liability rules are tractable for any possible combination of valuation distributions, the present-day judiciary will even less likely implement these rules accurately than the variants mentioned above.

B.   *Selecting the Distribution*

Continuing bravely on, we find that if a court could determine the allocatively optimal damages, it could again vary the litigants' individual expected payoffs without affecting the litigants' allocative choice. In other words, we can extend our previous convexity result to include second-order liability rules.

---

114. Formally, the density functions would take the form: $f(v) = 6v(1-v)$.

115. An example with the much simpler uniform distribution is derived in Ayres & Balkin, *supra* note 9.

Consider, for instance, the familiar example of two litigants with uniformly distributed valuations between $0 and $100. As discussed earlier, one example of a second-order rule with optimal damages would give the defendant a first-stage taking call option with a $44.44 exercise price and the plaintiff a second-stage take-back call option with a $66.66 exercise price. These nominal damages would induce the optimal pivot strategies in that the defendant would take when its value is greater than $33.33 and the plaintiff will take back when its value is greater than $66.66.[116]

But the same pivot strategies can be induced by a regime that gives the defendant the first-stage choice of paying $24.44 to the plaintiff to take the entitlement (D_c= $24.44) or receiving $20 from the plaintiff if it does not take (D_{NC} = -$20), and giving the plaintiff the same second-stage choice of paying $66.66 to take back the entitlement (D_C = $66.66; D_{NC} = $0). Under such a rule, the defendant will still perceive a $44.44 cost to taking in the first stage, and thus will continue to find it worthwhile to strategically take whenever its value is greater than $33.33. But again, by changing the first-stage option into a pay-or-be-paid rule, this implementation effectively shifts some of the expected total payoff from the plaintiff to the defendant. All that matters allocatively is the equality between the difference between D_C and D_{NC} and the nominal damages that would be optimal under a more traditional call or put rule. As emphasized above, calculating this initial number becomes computationally awkward, but a court can easily move D_C and D_{NC} in lockstep to create pay-or-be-paid or pay-or-pay rules that produce distributively different, but allocatively identical, equilibria. Even in the implementation of second-order rules, a court can make its distributive decision independently of its allocative decision. So just as there is an infinite class of rules that implement the two single-chooser allocations and an infinite number of rules that implement the two dual-chooser allocations, there turns out to be an infinite class of rules that can implement the two second-order allocations.

C. *Selecting the Chooser*

Which of these two second-order equilibria will be allocatively superior — the "unconstrained defendant-presumption" allocation or the "unconstrained plaintiff-presumption" allocation? The answer

---

116. Notice how this equilibrium comports with the dispositive takings principle. The defendant takes only when the defendant values the entitlement more than those plaintiffs who will not take it back. The class of plaintiffs who will not take back has valuations ranging from $0 to $66.66, and the defendant will take when valuing the entitlement more than the mid-point of this range, $33.33. Similarly, the plaintiff will take back only when the plaintiff values the entitlement more than those defendants who took it in the first stage. The class of defendants from whom a plaintiff might take back has valuations ranging from $33.33 to $100, and the plaintiff will take back only when valuing the entitlement more than the midpoint of this range, $66.66.

parallels our earlier analysis of the constrained dual-chooser alloca-
tions: it depends on the litigants' relative mean valuations. The over-
arching principle continues to be that as between the two second-
order rules, courts should tend to allow either party to veto allocations
to the party with the lower mean valuation. This means that the un-
constrained plaintiff-presumption equilibrium will tend to operate
more efficiently than the unconstrained defendant-presumption equi-
librium when

$$\mu_\Pi > \mu_\Delta.$$

For example, imagine that the defendant's valuation is uniformly
distributed between $0 and $100 and that the plaintiff's valuation is
uniformly distributed between $80 and $180. We can show that the op-
timal pivot values for an unconstrained plaintiff-presumption alloca-
tion equal $86.66 and $93.33 for the defendant and plaintiff, respec-
tively.[117] In contrast, the optimal pivot values for an unconstrained
defendant-presumption allocation equal $60 for the plaintiff and $120
for the defendant, respectively.[118] Each of these pivot values lies out-
side the range of the litigant's valuation distributions, and, as a result,
an unconstrained defendant-presumption allocation with optimal
damages would always allocate the entitlement to the plaintiff. Figure
9 graphically depicts the two allocations induced by the two rules, with
optimal pivot values.[119] As shown in the figure, the unconstrained
plaintiff-presumption rule is more efficient than an unconstrained
defendant-presumption rule (as the relative difference in misallocated
entitlements demonstrates). The unconstrained plaintiff-presumption
rule on balance does a better job of allocating the entitlement to the
defendant in the unusual circumstance when the defendant has the
higher value, whereas the unconstrained defendant-presumption rule
can do no better than allocating the entitlement solely to the plaintiff.

---

117. By the dispositive takings principle, the pivot values must simultaneously satisfy the
following two equations:

$$Pivot_\Delta = \frac{80 + Pivot_\Pi}{2}, \text{ and } Pivot_\Pi = \frac{Pivot_\Delta + 100}{2},$$

which, when solved, yield the values in the text. The nominal damages that would induce
these pivot taking values equal $92.43 for the defendant's takings and $93.33 for the plaintiff
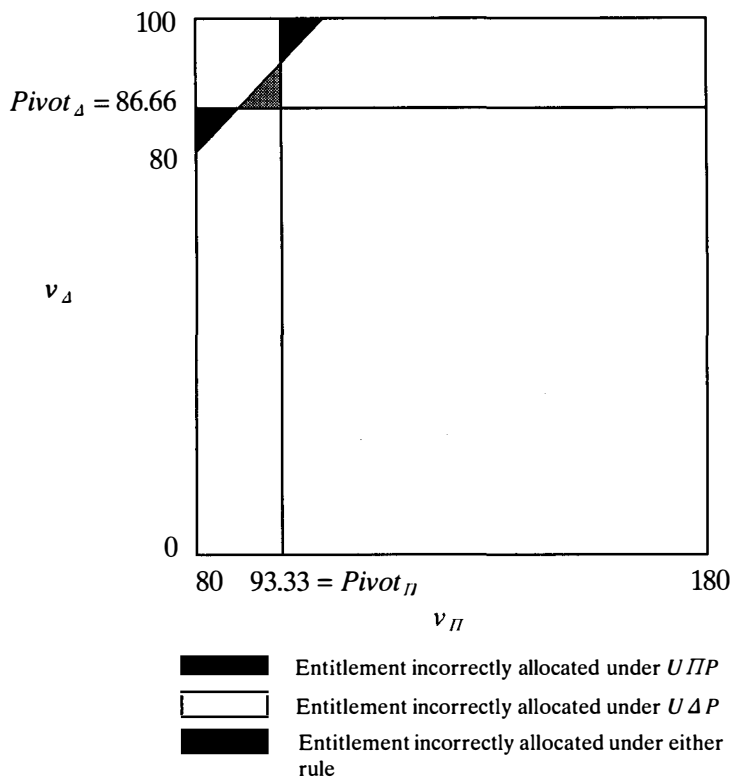taking back.

118. According to the dispositive-takings principle, the pivot value must simultaneously
satisfy the following two equations:

$$Pivot_\Delta = \frac{180 + Pivot_\Pi}{2}, \text{ and } Pivot_\Pi = \frac{Pivot_\Delta}{2},$$

which, when solved yield the same values in the text.

119. *See* Ayres & Balkin, *supra* note 9, at 727 fig.1.

Figure 9

Entitlement incorrectly allocated under $U \Pi P$

Entitlement incorrectly allocated under $U \Delta P$

Entitlement incorrectly allocated under either rule

D. *Selecting Among the Six Foundational Allocations*

Finally, we can usefully compare the relative efficiency of all six fundamental allocations. If courts could accurately calculate optimal damages for second-order allocations, the better second-order allocation would always weakly dominate all other allocations in the sense of producing at least as high an expected total payoff for the litigants as any of the other five allocations. This is intuitively obvious because the second-order rules have fewer constraints than either single-chooser or dual-chooser rules. The court sets not one, but two allocative prices and thus has more freedom to harness the private information of both litigants. Setting two (or more) allocative prices allows a court to establish an internal auction — where each taking at a sequentially higher price (or, as is the case with put options, lower price) effectively allows the individual litigants to credibly signal more in-

formation about their private valuations.[120] Heuristically, these less constrained liability rules permit the court to make more nuanced assignments of allocative authority. Just as the convexity result permits courts to distribute the expected payoffs as they see fit, the higher-order liability rules permit courts to distribute allocative power as they see fit — instead of giving one side or the other the sole allocative authority (as with single-chooser rules), or giving the two parties symmetric allocative authority (as with the constrained dual-chooser rules).

More prosaically, second-order rules can never operate less efficiently than the other rules because, with the appropriate choice of damages, second-order rules themselves can effect the same allocations as either of the single-chooser or dual-chooser allocations. Courts can shape second-order rules to effectuate single-chooser allocations by setting the first-stage exercise price at the nonchooser's valuation and then setting the second-stage price at a sufficiently prohibitive amount so that the second-stage chooser would not exercise it. In addition, courts can shape second-order rules to effectuate dual-chooser allocations by simply setting the first- and second-stage exercise prices equal to one another.[121]

Nevertheless, much as we are attracted to the elegant way in which second-order rules harness the information of both litigants, even when one litigant has relatively poor information, we believe there are many pragmatic drawbacks to implementing second-order rules that rightly refocus our attention on the single-price dual-chooser allocations. This latter type of dual-chooser allocation often represents the best pragmatic alternative available to courts for harnessing both litigants' private information in allocating an entitlement.[122] The constrained dual-chooser allocations, by definition, require the courts to calculate only a single allocative price, which, in the normal case, will approach something close to the average of the litigants' mean valuations. The single-price dual-chooser allocations analyzed above also induce transparent, dominant strategies that both the court and the litigants can readily understand. If the court sets dual-chooser damages at $\$D_{\Pi P}$ under a plaintiff-presumption allocation, the entitlement will go to the defendant only if $v_\Pi < D_{\Pi P} < v_A$. The relative simplicity and transparency of such allocations suggest that the theoretical

---

120. *See id.* at 729.

121. Indeed, Ayres and Balkin have shown that in a world with costless takings, successively higher-order rules can achieve successively higher expected total payoffs. *See id.* at 727-33. Higher-order liability rules create more refined auctions with smaller minimum bid increments, and as a theoretical matter, these rules in the limit can produce first-best efficiency.

122. Ayres and Balkin, *id.* at 733-36, have shown, for example, that as the transaction costs of taking or of calculating damages increase, single-price rules will tend to dominate multiple-allocative-price rules.

dominance of second-order regimes does not moot our initial efforts to identify how courts should select from an admittedly constrained class of single-price allocations.

IV. EXTENDING THE ANALYSIS

In this Section, we relax many of the more stringent assumptions of the foregoing analysis and assess the extent to which these assumptions limit the model's applicability or its normative implications.

A. *Administrative Costs*

The foregoing analysis ignored the costs of administering the liability rules. These costs will likely drive efficiency-minded lawmakers toward simpler rules. As just mentioned, second-order rules may, for example, entail high transactional or computational costs.[123] Indeed, when bargaining becomes impossible, the costs of liability rule administration may militate toward the use of property rules.[124] Property rules tend to deter nonconsensual takings and thus avoid the costs of administering a liability rule regime. When transaction costs remain relatively low, on the other hand, the litigation costs involved in administering a liability rule can somewhat counterintuitively facilitate greater allocative efficiency.[125] Under a liability regime, litigation costs give the parties an additional impetus to negotiate and hence can make liability rules more efficient than property rules.

Carol Rose has also stressed that courts may need to incur nontrivial costs in determining which parties should have a potential claim to ultimate ownership of particular entitlements.[126] The costs of determining who owns what entitlements will again militate in favor of property rules but should not greatly impact which of the liability rules will work best. The foregoing analysis suggests that optimal tailoring of a liability regime will turn on an assessment of the litigants' valuation distributions. We will return to these issues again when we consider the associated problems of "numerosity."

---

123. *Id.* at 717.

124. One still faces, however, the administration cost of determining who should be assigned the property right. Under certain circumstances, the costs of determining the high valuing litigant may exceed the costs of calculating damages. Richard R. W. Brooks, Choosing Property Rules or Liability Rules: The Burden of Determining Optimal Remedies (Aug. 2001) (unpublished working paper) (on file with authors).

125. Ayres and Talley, *supra* note 9, at 1065-72, stressed this point.

126. *See* Rose, *supra* note 9, at 2190-91.

B. Alternative Informational Assumptions

The foregoing analysis made imperfect information assumptions that are standard in the law and economics literature — namely, we assumed that the probability distributions of the litigants are (after trial) common knowledge, but only litigants know their own true valuations.[127] Here, we consider three variations on these assumptions: (a) the litigants are imperfectly informed about their own true valuations; (b) the litigants' probability functions are not common knowledge; and (c) courts systematically misestimate the probability distributions of the litigants' values.

The possibility that litigants do not precisely know their own valuations does not seriously limit our analysis. Although litigants may not know with precision the prospective value of owning an entitlement, the litigants nonetheless might estimate a certain cash-equivalent amount that would leave them indifferent as between the entitlement or this cash amount. Just as individuals often buy and sell entitlements for discrete cash amounts despite their uncertainty about the future value of ownership, litigants may be able to estimate their own values.

As long as the individual litigants are better informed than the court, liability rules — by privatizing the ultimate allocative authority — can enhance ex post efficiency. In mathematical terms, as long as litigants' beliefs about their valuation probability functions have less volatility than the court's beliefs about the litigants' valuations, the litigants will have an informational advantage over the court in ultimately allocating the entitlement.[128]

Even if litigants know their own valuations, one might worry that they cannot, contrary to our model's assumption, assess even the probability distribution of the other party. Nonetheless, we believe that this should also not pose a large concern. Liability rules do not require a chooser to assess directly the probability functions of an adversary; rather, a chooser need only confront the court-determined damage amount, or what we have called the allocative price. Court-appointed choosers do not need to estimate nonchooser valuations as long as the court itself, after discovery and trial, can estimate a valuation distribution sufficiently to determine optimal single- or dual-chooser damages.[129] As stressed above, single-chooser allocations will

---

127. *See, e.g.*, Kalyan Chatterjee & William Samuelson, *Bargaining Under Incomplete Information*, 31 OPERATIONS RES. 835 (1983); Peter C. Cramton, *Strategic Delay in Bargaining with Two-Sided Uncertainty*, 59 REV. ECON. STUD. 205 (1992).

128. *See supra* note 53 and accompanying text.

129. In some contexts, especially with regard to dual-chooser regimes, a potential litigant makes the first-stage choice before the court has determined damages. For example, a potential defendant may decide to encroach in anticipation of the plaintiff having the put-like choice of injunction or damages for permanent encroachment. In such circumstances, a rational encroacher would want to assess what damages the court will likely set.

also tend to simplify the court's decisionmaking by requiring the court to estimate only the mean of the litigant with the less speculative valuation. Of course, the less speculative valuation may still, from the court's perspective, remain quite speculative, but courts only need to make an unbiased estimate in order to harness the litigant's private information in allocating the entitlement.

A more serious informational concern involves court estimates of litigants' valuation distributions that for some reason become systematically biased. For example, if courts in implementing single-chooser allocations systematically underestimated (overestimated) the non-chooser means, then they would set damages too low (high) and would induce the chooser to inefficiently allocate the entitlement too often (rarely) to itself. If higher branches of government (appellate courts or the legislature) realize that lower branches (trial courts/juries) systematically underestimate the mean value of residents, a standard response is that the higher branch might force the lower branch to eschew liability rules altogether and instead give the residents the pollution entitlement protected by a property rule. If the mean valuation assessments of the more authoritative branches of government are also biased, the allocation problem becomes truly intractable. We readily admit that our model will not perform well, but neither will most other allocative regimes.

C. *Unintentional Taking*

Our option model has been quintessentially about *intentional* takings. One or both litigants choose whether to exercise intentionally a particular option. In many contexts, this assumption is reasonable. Potential defendants at times make deliberate decisions to impair plaintiffs' entitlements. The plaintiffs' election of remedies to seek an injunction or monetary damages also almost always fits the deliberate model.

Many, if not most, torts, however, concern unintentional takings. We therefore must ask whether our models have anything to say to this large class of takings. We believe they do. First, unintentional takings actually increase the ability of courts to apply single-chooser allocations, especially plaintiff-choice allocations. If Calabresi inadvertently and nonstrategically encroaches on Melamed's land, courts have even greater freedom to apply a true plaintiff-choice allocation and give Melamed unilateral allocative authority based on a single allocative price. With regard to deliberate encroachments, by contrast, courts may not be able to apply a pure plaintiff-choice allocation, because potential encroachers may take into account the second-stage plaintiff choice when deciding whether or not in the first stage to encroach.

Second, and more speculatively, an options model of deliberate taking may have something to tell us about negligence regimes. Even when potential defendants do not deliberately take plaintiffs' entitlements, the defendants often do make deliberate choices as to the care level they will use to avoid taking. In some contexts we can think about choice of care level as tantamount to choosing the probability that an accident will occur. Put more in the argot of finance, level of care choices represent the potential defendants' choice of probability that they will exercise their takings options. Damages that equal the nonchooser's mean valuation — that is, the plaintiff's mean valuation of impairment — will induce takings only when the probabilistic value of takings, as measured in saved due care costs, exceeds the plaintiff's average loss. Thus, we can restate the Learned Hand formula in option terms, but we often do not gain much from the translation. Many torts do not fit a plaintiff-choice allocation because once the defendants destroy the entitlements, a court cannot ask the plaintiffs whether they would prefer the entitlements themselves (for example, an arm) instead of cash.

Even with regard to negligent taking, however, some circumstances may permit a second-stage plaintiff election of remedy, and the foregoing analysis may apply. If Epstein negligently takes Calabresi's galoshes at the theater, common law courts can give, and traditionally have given, Calabresi the option to take back the shoes or to receive money.[130] In setting the damage amount, the court may create a kind of dual-chooser allocation that would impact both the due care with which Epstein examines galoshes at the theater and Calabresi's decision whether to seek injunctive or damage relief.

D. *Numerosity*

Our previous discussion has relentlessly focused on bilateral disputes. We have assumed that a dispute concerns only one plaintiff and one defendant; and we have also assumed that there are no third parties who might contend for being the highest valuer of the entitlement. Both of these assumptions characterize a nontrivial amount of real-world litigation. For example, in *Copart Industries v. Consolidated Edison*, a particular type of pollution from a ConEd plant disrupted the ability of a single plaintiff to continue its new car preparation business.[131] Many nuisance disputes, however, will involve multiple litigants on one or both sides of the litigation. For example, the tendency of pollution to spread means that a single defendant's effluence

---

130. *See* Ayres, *supra* note 9, at 794.

131. 362 N.E.2d 968 (N.Y. 1977); *see also* cases cited *supra* note 27.

may affect multiple plaintiffs.[132] Krier and Schwab have seized on this difference in numerosity to argue that defendants will tend to make more efficient choosers because they are less numerous.[133] As applied to our analysis, Krier and Schwab would argue that defendant-choice rules, including the Rule 2 call-option implementation and the Rule 5 put-option implementation, would tend to dominate plaintiff-choice rules. We agree that the additional collective action difficulties of a diffuse plaintiff class will, all else being equal, tend to militate toward the use of defendant-choice allocations. Furthermore, our decoupling result improves upon Krier and Schwab's theory by showing that a continuum of defendant-choice allocations can take equivalent advantage of the defendant's private information while providing even more nuanced distributive possibilities.

Nevertheless, we disagree with Krier and Schwab's seemingly exclusive focus on numerosity as the determinant of the more efficient chooser. Instead, we believe that selection of the more efficient chooser will also turn importantly on which party has a greater informational advantage relative to the court. All else equal, the party the court believes has the more volatile valuation will be the more efficient chooser.

In many nuisance disputes numerosity and volatility will push the court in different directions. The plaintiffs in nuisance disputes often include more members than the defendants, but these plaintiffs are also likely to have more speculative valuations. Often, courts may more easily measure the economic costs of abating or discontinuing pollution than assess the noneconomic damages of pollution on surrounding Residents. Even though multiple Residents will often have high collective decisionmaking costs, they will also often possess more relevant private information in that they will know how much pollution really harms them.[134]

At times, courts should give the choice to multiple Residents and accept the resultant collective-action problems, rather than give the Polluter the choice and disregard the Residents' private information. For example, we believe that the residents in *Boomer v. Atlantic Cement*[135] might have made better choices than the polluter, notwithstanding their relative numerosity. Relative to the court, these residents would have had much more information about their own valuations of the entitlement. While we concede that numerosity will

---

132. *See* Richard L. Revesz, *Federalism and Interstate Environmental Externalities*, 144 U. PA. L. REV. 2341, 2342-44 (1996).

133. Krier & Schwab, *supra* note 7, at 471-72.

134. In addition, Residents often will have higher levels of risk aversion than the Polluter, who may in turn have greater opportunities to diversify loss, and hence, Residents will, other things being equal, tend to make better choosers.

135. 257 N.E.2d 870 (N.Y. 1970).

often militate against giving the choice to multiple residents, the variance effect will often militate in favor of precisely this brand of allocative choice.

Our two-player models also limit the applicability of our analysis with regard to the exclusion of potential third-party claimants for ultimate ownership of the entitlement in question. In a world in which three or more different parties are the potentially most efficient ultimate owner of the entitlement, courts will have more difficulty implementing a liability rule regime. Absent bargaining, however, property rules will also perform badly in allocating the entitlement to the most efficient owner, unless the imperfectly informed court by chance happens to allocate the entitlement to its highest valuer. When the litigants can feasibly bargain, the multiplicity of potential highest valuers favors a property rule. Liability rules, by contrast, can increase the cost of contracting (a) by failing to allocate claims to its potentially highest valuers, or (b) by allocating claims to people who do not value it highly. For example, when a court gives a partial entitlement to inefficient owners, they will only turn around and try to sell it, and their usual incentives to seek inflated prices can impede negotiations with the class of potentially highest valuers. Thus, low valuers who mistakenly receive a claim on the entitlement can inefficiently hold up the transaction.[136] This is also another example of what Heller called the anticommons problem: too many people can block the transfer of an entitlement to its highest valuing owner.[137] As the pool of potentially high valuers increases, liability rules of both the single-chooser and dual-chooser varieties can exacerbate this holdout or anticommons problem.

E.   *Cognitive Bias and Wealth Effects*

A central finding of the previous analysis concerns the allocative invariance of all implementations within a particular fundamental allocative class. For example, we have shown that any defendant-choice rule with a particular allocative price (D_A) will produce identical allocations for any possible sets of private plaintiff and defendant valuations. This invariance result only holds true, however, in the absence of wealth and cognitive framing effects.

As with the Coase theorem's invariance prediction,[138] our allocative invariance finding will not hold true if the chooser's willingness to

---

136. *See* Ayres & Talley, *supra* note 9, at 1086.

137. Heller, *supra* note 100.

138. *See* John J. Donohue, III, *Diverting the Coasean River: Incentive Schemes to Reduce Unemployment Spells*, 99 YALE L.J. 549, 550 (1989); Stewart Schwab, *A Coasean Experiment on Contract Presumptions*, 17 J. LEGAL STUD. 237, 242-43 (1988).

pay is affected by its overall wealth.[139] Defendant-choice rules that allocate larger expected payoffs to the defendant might make the defendant more or less willing to pay a particular allocative price in allocating the entitlement to itself.

More generally, the extent to which choosers' willingness to pay a given price differs from their willingness to accept that same price qualifies the allocative invariance result. For both cognitive- and wealth-effect reasons, particular choosers will often demand higher prices when selling an entitlement than they would willingly pay to buy it.[140] As Saul Levmore has noted in discussing put options:

> Armed with the legal right to be free of pollution, [a Resident] might have demanded $1,000,000, for example before allowing [a Polluter] to send effluents her way, so that [the Resident's] damages from [the Polluter's] operation can fairly be said to amount to $1,000,000; but this does not mean that [the Resident] could or would pay this amount to stop [the Polluter], if the law requires [the Resident] to pay in order to [stop pollution].[141]

This suggests that a range of prices exists (in Levmore's example, between $700,000 and $1,000,000) in which the plaintiff-residents would allocate the entitlement to themselves under a put-option implementation (Rule 6) but would allocate the entitlement to the defendant under a call-option implementation (Rule 4). For these intermediate damage amounts, the law's selection of an initial entitlement holder not surprisingly will tend to be sticky. This should not be surprising. Scholars have long known that endowment effects impede trade and bargaining in the shadow of property rules; this result carries over to option regimes as well.

More controversial, however, is whether lawmakers should privilege endowment effects in a cost-benefit calculus. If Residents value an entitlement more than a Polluter merely because the court initially endows them with it, lawmakers might want to set a low call-option price to allow the Polluter to take control of the entitlement and benefit from its own endowment effect. One goal, which courts must trade off against other equity goals, focuses on production of a legal regime where the entitlement flowed to the person who has the highest will-

---

139. *See* A. MITCHELL POLINSKY, AN INTRODUCTION TO LAW AND ECONOMICS 7-10, 119-27 (2nd ed. 1989); Stewart J. Schwab, *Coase Defends Coase: Why Lawyers Listen and Economists Do Not*, 87 MICH. L. REV. 1171, 1178-83 (1989) (book review).

140. Mark Kelman, *Consumption Theory, Production Theory and Ideology in the Coase Theorem*, 52 S. CAL. L. REV. 669, 673 (1979); Duncan Kennedy, *Cost-Benefit Analysis of Entitlement Problems: A Critique*, 33 STAN. L REV. 387, 401 (1981); Russell Korobkin, Note, *Policymaking and the Offer/Asking Price Gap: Toward a Theory of Efficient Entitlement Allocation*, 46 STAN. L. REV. 663 (1994).

141. Levmore, *supra* note 9, at 2166.

ingness to accept.[142] From an efficiency standpoint, this goal would rank people by their value of "having" rather than their value of "not having." Other things being equal, this standard would lead toward giving the initial entitlement to the party with the larger endowment effect.

The proliferation of different single-chooser rules, including the call, put, pay-or-be-paid, and pay-or-pay implementations, again, can help lawmakers accommodate competing concerns. When cognitive or wealth effects cause these implementations to differ allocatively, this proliferation will at least permit policymakers more flexibility in balancing a variety of efficiency and equity concerns. Efficiency might demand that the initial entitlement go to the party with the larger endowment effect. But if this is the party that the court believes also is the more efficient chooser, then the court can accommodate both of these efficiency concerns by giving the party both the initial entitlement and a put option.

But more work needs to be done to see the extent to which an endowment effect persists under the various implementations of single- and dual-chooser allocations. The path-breaking mug experiments of Kahneman, Knetsh, and Thaler only tested whether entitlements protected by property rules gave rise to endowment effects.[143] It is much less clear whether endowments protected by liability rules, that is, subject to someone else's call option, give rise to the same effect. In the shadow of a call option, the initial entitlement holder has a more transient, less secure claim to the entitlement and thus might not establish the cognitive bond that gives rise to the effect. In fact, Jeff Rachlinski and Forest Jourden have recently conducted experiments showing that traditional liability (call-option) rules largely extinguish the endowment effect.[144] But put options or other intermediate implementations, such as the pay-or-be-paid rule may produce different

---

142. *See* Herbert Hovenkamp, *Legal Policy and the Endowment Effect*, 20 J. LEGAL STUD. 225, 232-34, 245-47 (1991) (claiming that willingness to accept should always be used rather than willingness to pay). Russell Korobkin, however, has suggested reasons why allocation of entitlements on the basis of willingness to accept is inappropriate when independent social norms oppose sales of certain types of entitlements:

> Narrowly stated, our society teaches us that we should not sell some things, although buying them is socially acceptable. Although an individual might, in a cultural vacuum, prefer a specific quantity of money to a certain entitlement, he might feel that selling the entitlement is improper and therefore, his WTA will exceed his WTP. This explanation has been called the "dignity hypothesis."

Korobkin, *supra* note 140, at 691-92.

143. Daniel Kahneman et al., *Experimental Tests of the Endowment Effect and the Coase Theorem*, 98 J. POL. ECON. 1325 (1990).

144. Jeffrey J. Rachlinski & Forest Jourden, *Remedies and the Psychology of Ownership*, 51 VAND. L. REV. 1541, 1545 (1998) (arguing that "the endowment effect might depend on whether the law protects an interest with a property rule or a liability rule"). This possibility was first suggested in Ayres & Talley, *supra* note 9, at 1102.

endowment effects. For example, put-option rules may heighten the endowment effect because put-option implementations give the initial entitlement holder even more control over the "endowment" than traditional property rules, and thus may enhance the solidity of owner-ship.[145] If we view the endowment effect as a source of real value that we want to nurture, as opposed to a "status quo bias" or "framing effect" that we want to overcome,[146] then put options may overcome the inefficiencies of the holdout problem without destroying the "value of having." More generally, however, the infinity of single-chooser rules discovered here gives policymakers more flexibility in testing for and managing the pros and cons of cognitive bias.

F.   *Correlated Valuations*

Our analysis thus far has assumed that the litigants have independently distributed valuations. While we can plausibly assume this in some contexts,[147] Kaplow and Shavell have pointed out that a plaintiff's and defendant's valuation of physical objects may correlate.[148] If a plaintiff places a high value on an oilfield or laptop computer, there are good reasons to think that a defendant's value may also be high.

Correlated valuations make it more difficult for courts to use liability rules to harness the parties' private information. For example, Kaplow and Shavell have shown that if we can decompose the litigants' overall valuations into a "common-value" and an "idiosyncratic-value" component,[149] courts will have difficulty in setting a damage amount that induces takings so as to increase social welfare. For example, suppose that (i) the litigants have a uniformly distributed common-value component between $0 and $100; (ii) the plaintiff has a uniformly and independently distributed idiosyncratic-value component between $0 and $100; and, (iii) the defendant has a

---

145. Russell Korobkin, however, has suggested to us that the put option may commodify the entitlement in the owner's mind in ways that reduce the endowment effect. Imagine that after a pure-bred dog has puppies, a mother tells her daughter she can keep the puppies or, after they are weaned, can sell them for $50 to a neighbor who has made a firm offer. The daughter's knowledge that she has a put option might stop her from bonding as closely with the puppies during the weaning period as she would have otherwise and hence mitigate the endowment effect.

146. *See* Russell Korobkin, *The Status Quo Bias and Contract Default Rules*, 83 CORNELL L. REV. 608 (1998).

147. For example, in many nuisance contexts it is plausible to assume that a Polluter's value from polluting will be independent of a Resident's value of nonpollution.

148. Kaplow & Shavell, *supra* note 15, at 759.

149. *Id.* at 759-60 ("First, suppose that things have a significant *common value*, that is, a component of value that is the same for both the owner and any taker.... Second, assume that things also have idiosyncratic value to individuals. *Idiosyncratic value* derives from characteristics of a thing that different individuals evaluate differently, such as the design of a home.").

uniformly and independently distributed idiosyncratic-value component between $0 and $75. A traditional liability rule, Rule 2, with damages set equal to the plaintiff's mean valuation of $100 (a $50 mean common value plus a $50 mean idiosyncratic value) will operate less efficiently than a property rule, Rule 1, which simply gave the plaintiff the entitlement.

When the litigants have positively correlated valuations, the court will have difficulty setting an allocative price to distinguish between unexpectedly high realizations of the common-value component versus unexpectedly high realizations of the idiosyncratic-value component. A Rule 2 with $100 damages would induce the defendant in this example to take too often. For example, when the common-value component turned out to equal $80, a defendant would take if the defendant's idiosyncratic value fell as low as $20, even though such takings would inefficiently allocate the entitlement because the defendant's mean idiosyncratic value equals $50.[150]

Restated in the terms of this Article, Kaplow and Shavell's suggestion means that when the litigants have substantially correlated valuations, the optimal delegation is no delegation: courts should allocate the entitlement via a property rule rather than delegate the allocative decision to the litigants via a liability rule.

Although Kaplow and Shavell are correct that correlated valuations make it more difficult to harness private valuations, they wrongly conclude that property rules work better than liability rules when litigants have positively correlated values. In a separate essay, we show that the examples in Kaplow and Shavell's original article systematically overstate the advantages of property rules by comparing the more efficient property rule to a liability rule that has nonoptimal damages and the less efficient chooser.[151] But for now, it is sufficient to see how nicely dual-chooser rules can respond to the correlated value problem.

Remember that as a traditional liability rule, Rule 2 has the problem that defendants may take not because they have high idiosyncratic values but merely because both litigants have unexpectedly high realized common-values. The court would like defendants to take only if their idiosyncratic values are unexpectedly high, but the court cannot observe either their common- or idiosyncratic-value components. It can only set a single damage figure that might induce defendants to take either because they have unexpectedly high idiosyncratic values (creating efficient takings), or because both they and the plaintiffs have an unexpectedly high common value (creating inefficient tak-

---

150. *Id.*, at 760.

151. *See* Ian Ayres & Paul M. Goldbart, A Critique of "Tangibility" as the Basis for Property Rules (2001) (unpublished manuscript) (on file with authors).

ings). Although the court cannot observe whether defendants take because of a high common value or a high idiosyncratic value, the plaintiff can.

Dual-chooser rules respond well to the problem of correlated valuations because these rules allow the other side to *veto* takings driven by common-value realizations. The problem with Rule 2 in the foregoing example was that defendants might take merely because the common-value component was unexpectedly high. A plaintiff-presumption allocation, however, eliminates this problem by allowing the plaintiff to veto takings that are driven by both parties having a high common value. Under a plaintiff-presumption allocation, the entitlement will ultimately go to the defendant only if the plaintiff's total value, common plus idiosyncratic, is less than the damage amount and only if the defendant's total value is greater than the damage amount. Indeed, in the foregoing example we can show that the optimal dual-chooser rule produces systematically higher expected joint payoffs than the optimal property rule.[152] Thus, while correlated valuations make optimal delegation a more difficult judicial task, one or two private heads can still be better than none.

G. *Bargaining*

Probably the most serious qualification to our analysis thus far concerns our assumption of autarky. Our formal model assumes that the litigants cannot bargain with each other and considers instead the equilibrium allocation and distribution absent negotiations. This assumption can be given two partial defenses, but ultimately much more work needs to be done to understand how liability rules should be structured in a world where transaction costs are relatively small. In this Section, we sketch these two partial defenses for our autarky assumption and then show how dual-chooser allocations induce a new type of information forcing that may favor their use when transaction costs are relatively low.

The first partial defense of an autarky assumption is empirical. In at least some bilateral contexts, the disputants seem extremely reluctant even to negotiate, much less agree to contract. Ward Farnsworth, for example, surveyed nuisance disputants after an initial trial decision and found almost without exception that the litigants and their attorneys failed even to consider negotiation in the shadow of the law.[153] Enmity between the litigants may produce an important barrier to trade and thus inhibit private contract as a means of allocation.

---

152. In the foregoing example, Rule 1 produces an expected payoff of $100, while the optimal plaintiff-presumption dual-chooser rule produces an expected payoff of $101.39.

153. Farnsworth, *supra* note 29.

Second, an understanding of how liability rules perform when litigants cannot feasibly trade may provide valuable insights into the relative efficiency of these rules when the litigants can bargain. The autarkic payoffs that we have analyzed above become the threat points that will deeply influence the litigants' bargaining strategies. Negotiation theorists at times refer to such autarkic payoffs as representing the individual litigants' "reservation prices" for trade, or their "BATNAs" ("best alternative to negotiated agreement").[154] Because litigants cannot be forced through consensual trade to accept less than their BATNA, our previous analysis of how judges might vary the distribution of payoffs will continue to hold true in a world where bargaining feasibly can occur. For example, adopting a defendant-choice allocation with a relatively high alpha value (a pay-or-be-paid rule) enables the court, even when bargaining *can occur*, to assure the plaintiff a higher expected return and the defendant a lower expected return than a defendant-choice rule with a put option (Rule 5).

Kaplow and Shavell, however, have gone further to argue that legal regimes that are relatively more efficient in allocating the entitlement when bargaining is not possible are likely to continue to be more efficient as transaction costs subside and bargaining becomes feasible. Of course, the Coase theorem teaches that all contractable legal regimes, such as property, single-chooser, and dual-chooser rules, will allocate entitlements equally efficiently when transaction costs are nil. Nevertheless, Kaplow and Shavell have argued that the relative efficiency of different rules when bargaining cannot take place will tend to *persist* when bargaining becomes possible, but not perfectly.[155] Under this persistence conjecture, the nonconsensual advantage that, say, a defendant-choice allocation might have (when the litigants' mean valuations are similar and the variance of the defendant's valuation is greater than the plaintiff's) relative to the other three single-price allocations would be likely to persist if the parties could bargain. For Kaplow and Shavell, this nonconsensual headstart that a particular rule has over others in promoting allocative efficiency will on average persist as transaction costs recede and trade becomes possible.

If the persistence conjecture is true, then the foregoing analysis would end the story. By knowing how best to tailor rules when bargaining was prohibited, judges would perforce know how best to tailor rules when bargaining was allowed. Unfortunately, we are not convinced that the persistence conjecture tells the entire story. While it is surely the case that the most efficient liability rule under autarky has a headstart when bargaining is permitted, it is at least possible that other,

---

154. *See* Jennifer Gerarda Brown & Ian Ayres, *Economic Rationales for Mediation*, 80 VA. L. REV. 323, 331 & nn.26 & 27 (1994).

155. Kaplow & Shavell, *supra* note 15, *passim*.

consensual attributes of the rules may come to dominate the noncon-
sensual attributes. We know, for example, that the persistence conjec-
ture at most describes only a tendency, because Ayres and
Talley have provided a numerical counterexample in which a noncon-
sensual advantage of property rules over liability rules did not persist
once bargaining became possible.[156] Moreover, these authors have
demonstrated that liability rules have a much stronger "information-
forcing" effect than property rules do. This information-forcing effect
represents a potential "consensual advantage" of liability rules that
may be more important in the relative efficiency of different legal re-
gimes when bargaining costs are lower than the nonconsensual advan-
tage that Kaplow and Shavell have emphasized.

In terms of the foregoing analysis, single-chooser allocations give
nonchoosers a robust incentive to reveal through bargaining whether
their value is above or below their mean valuation. For example, imag-
ine that a court establishes a defendant-choice call-option rule in
which defendants can self-allocate the entitlement when they make a
payment to plaintiffs of one million dollars, the plaintiffs' estimated
mean valuation. The plaintiffs can respond by trying either (a) to sell
the entitlement for less than one million dollars, or (b) to buy the de-
fendants' call option (in essence trying to bribe the defendants not to
take nonconsensually). The plaintiffs as nonchoosers will offer to pay
the defendants not to take only if the plaintiffs have an actual valua-
tion higher than the court's estimate of their mean valuation, and they
will offer to sell the underlying entitlement only if they have an actual
valuation lower than the court's estimate. By merely expressing an in-
terest in one type of trade or another, nonchoosers can thus credibly
signal whether their actual valuation is higher or lower than the court's
estimate of their mean valuation.[157] Analogously, plaintiff-choice allo-
cations give the defendants as nonchoosers a robust incentive to signal
credibly how their actual valuation compares to their mean valuation,
again, merely by which type of trade they offer to transact.

If the parties' imperfect information poses the primary barrier to
trade, it is at least possible that these information-forcing effects will
dominate the nonconsensual advantages of the different liability rules
outlined above. In particular, the more efficient choosers in the ab-
sence of bargaining will clearly tend to be the litigants with the more
speculative valuation. When bargaining can occur, however, the more
efficient choosers might be the litigants with the less speculative valua-

---

156. Ayres & Talley, *supra* note 9, at 1104.

157. Under a put implementation of a defendant-choice rule, the plaintiffs would offer
to pay the defendants not to exercise their put only when the plaintiffs have an actual valua-
tion lower than the court's estimate of their mean valuation, and the plaintiffs would offer to
pay the defendants for the underlying entitlement only when the plaintiffs have an actual
valuation higher than the court's estimate of it.

tion, for negotiations will quickly allow high variance litigants to narrow their observable distributions by credibly signaling whether they have a relatively high or low valuation. To date, however, we have not succeeded in modeling an example in which, when bargaining can occur, the litigant with the less speculative valuation became the more efficient chooser. But, at a minimum, the information-forcing effect on the nonchooser systematically narrows the allocative efficiency differential between the two single-choice allocations.

The question of whether our autarkic results persist in the presence of bargaining becomes all the more interesting when we add to the mix the possibility of dual-chooser allocations. Dual-chooser allocations turn out to have an arguably even stronger information-forcing effect than single-chooser allocations. While single-chooser allocations give one of the litigants, the nonchoosers, an incentive to signal whether they have a relatively high or low valuation, dual-chooser allocations give both litigants an opportunity to signal credibly whether they have a high or a low valuation by merely expressing an interest in a particular type of trade.

Consider, for example, a plaintiff-presumption allocation implemented through a call-call option regime in which the plaintiff receives (a) the initial entitlement, (b) subject to the defendant's call option (with an exercise price of $D_{\Pi P}$), but (c) also a call option (with an exercise price of $D_{\Pi P}$) to take back the entitlement. In the absence of bargaining, the plaintiff-presumption allocation gives rise to two different types of inefficiencies. The first type occurs when the defendant has the higher valuation but the plaintiff nonetheless vetoes a defendant allocation by exercising its take-back option. This occurs when $v_A > v_\Pi > D_{\Pi P}$. And the Coasean solution would be for the defendant to buy the plaintiff's take-back (call) option. The second type of inefficiency occurs when the defendant has the higher valuation but the defendant nonetheless vetoes a defendant-allocation by failing to exercise its initial call option. This occurs when $D_{\Pi P} > v_A > v_\Pi$. And the Coasean solution for this inefficiency would be for the defendant to buy the underlying entitlement (for some price less than $D_{\Pi P}$).

When negotiating in the shadow of this legal rule, the defendants can credibly signal their relatively high valuation ($v_A > D_{\Pi P}$) by merely expressing an interest in paying the plaintiff not to exercise its take-back option.[158] Conversely, the plaintiff could credibly signal that it has a relatively low value ($v_\Pi < D_{\Pi P}$) by merely expressing an interest in selling its entitlement for a price less than $D_{\Pi P}$. Low-value defendants would never have a strategic reason for expressing an interest in buying — much less actually offering to buy — the plaintiff's take-back

---

158. The defendant could also credibly signal its high value by exercising its first-stage call option. But the information-forcing stressed in the text allows the defendant to signal its high valuation even before it is called upon to exercise its first-stage option.

option. And high-value plaintiffs would never have a strategic reason for expressing an interest in selling — much less actually offering to sell — their entitlement for a price less than $D_{\Pi\dot{p}}$. Thus dual-chooser allocations create some opportunity for both litigants to credibly communicate their relative valuations through the process of negotiation.

Just as correlated litigant valuations in the discussion above expanded the set of circumstances in which dual-chooser allocations yielded greater allocative efficiency, the arguably stronger information-forcing properties of dual-chooser allocations may further militate for their use as trade in the shadow of these rules becomes more feasible. A full analysis of liability rules under conditions of low but nonzero transaction costs has eluded scholars. The particular ranking of the rules' relative efficiency might turn not just on the general characteristics of the litigants' valuation distributions but also on the specifics of the bargaining procedures that are assumed.

In this Section, we have tried to argue that autarkic results are of interest because (a) some bilateral disputants face high transaction costs; (b) the nonconsensual headstart that we have identified may tend to persist when bargaining is possible; and (c) the autarkic pay-offs that we have derived become the threat points that will significantly inform the bargaining behavior of the parties. Our finding that dual-chooser allocations create different and arguably stronger information-forcing effects than those previously identified for single-chooser allocations suggests, however, that the autarkic analysis may not tell the complete story. On this dimension we are forced to repeat the hackneyed and unsatisfactory conclusion that more research on this issue is warranted.

CONCLUSION

While we have often motivated our analysis with a nuisance example, this Article is really about a deeper and more canonical problem about how courts should allocate rivalrous entitlements when the court has poorer information than the litigants about who is the higher valuing owner. Most basically, we have considered a conflict in which the court — even after discovery and an adversarial trial — has only a probabilistic assessment of each litigant's value, whereas each litigant knows (or at least has better information) about how much it values the entitlement. In such a basic bilateral dispute with imperfect judicial information, the court faces a nontrivial problem of how best to allocate claims to the entitlement.

When the private litigants are better informed than the court, the court can usefully adopt an allocative rule that harnesses some of the litigants' private information. This Article has shown that courts may choose from four basic allocations. The traditional liability rules —

Rules 2 and 4 — exemplify two of these basic allocations: the plaintiff-choice and defendant-choice allocations. But we have shown that there are two other allocations — dual-chooser allocations — which not only have common-law analogs but which, under specific circumstances, can produce greater allocative efficiency than single-chooser allocations.

Moreover, we have shown there are an infinite number of implementations of each type of rule which are allocatively equivalent but which allow the court to distribute the total expected payoff of the litigants as it sees fit. This means that courts, in resolving bilateral disputes, can decouple allocative and distributive concerns. By selecting the more efficient chooser or choosers and by setting the efficient allocative price, a court can maximize ex post allocative efficiency. Moreover, independent of these core allocative decisions, the court has a free hand to distribute the expected payoffs as it wants. Thus, distributively concerned courts are free to pursue equity or promote ex ante efficiency without sacrificing their natural desire to see that the entitlement ends up in the hands of the highest valuing user (the ex post allocative efficiency concern). This decoupling result and our analysis of dual-chooser allocations are our central findings.

While we have focused on how courts should structure liability rules, our analysis does have implications for the equally important choice of when courts should opt for liability rules over property rules. Legal scholars have generated a number of rationales for why property rules might dominate liability rules.[159] We focus here on three rationales for property rules that the foregoing analysis qualifies. One of the oldest rationales states that property rules induce trade better than liability rules when transaction costs remain relatively low. Ayres and Talley seriously undermined this rationale by showing that liability rules have a greater information-forcing effect than property rules and by providing examples in which liability rules were more effective in inducing trade. And this Article has extended this criticism by showing that dual-chooser rules also have new and potentially stronger information-forcing qualities that may also give dual-chooser rules a consensual advantage over property rules.

As a second rationale, many authors, including Ayres and Talley,[160] have argued that property rules give individuals better ex ante investment incentives to create the disputed entitlement to begin with. The idea here is that Krier will be less likely to build a new house if Schwab has the option to take nonconsensually and pay Krier's mean valuation. This result makes sense, however, only if we limit our atten-

---

159. *See* Ayres & Talley, *supra* note 9; Epstein, *supra* note 9; Rose, *supra* note 9.

160. Ayres & Talley, *supra* note 9, at 1743.

tion to the call implementations of liability rules.[161] The aforementioned decoupling result shows that both single-chooser and dual-chooser implementations can yield even greater expected payoffs for either litigant than a property rule regime can yield in the absence of bargaining. The decoupling result seriously qualifies the investment rationale for using property rules because it suggests that courts can retain the ex post information-harnessing benefits of liability rules without undermining ex ante investment incentives.

Finally, and most recently, Kaplow and Shavell have argued that property rules will tend to work better than liability rules when the litigants have correlated valuations. They have asserted that liability rules will tend to be less efficient because choosers will allocate the entitlement to themselves not because they have high idiosyncratic valuations but because the common-value element of their valuations takes on an unexpectedly high value and exceeds the court's allocative price.[162] Their analysis makes eminent sense in comparing the efficiency of property rules to single-chooser rules,[163] but as we have demonstrated above, they did not consider the possibility of dual-choice rules, thereby limiting their analysis. Dual-chooser rules give both litigants a veto over certain allocations and can therefore better select for circumstances when the idiosyncratic and not the common valuation is driving the allocative choice.

In sum, the option analysis of liability rules as a theoretical matter tends to tilt increasingly toward the use of liability rules. In some sense this occurs as a natural byproduct of entering more liability-rule horses into the race to compete against the two property-rule stalwarts (Rules 1 and 3). Still, there are strong reasons why property rules will at times dominate,[164] but this Article and other recent efforts have made progress in identifying narrower but firmer grounds on which property rules can stake their superiority. When liability rules dominate, however, lawmakers would be well advised to consider the fuller

---

161. Even under the call implementation, it is not clear why this argument makes sense. If Krier at the time of investment does not yet know his true valuation, then promising him his mean valuation should not deter his investment. Indeed, if he is risk averse it may decrease the variance in his payoffs and make him more willing to invest. If, however, Krier at the time of investing knows his future valuation, then promising him his mean valuation, should Schwab take the entitlement, may deter him from some valuable types of investment when his actual value turns out to be higher than the court's perception of his mean valuation.

162. Conversely, choosers may allocate the entitlement to the nonchoosers, not because they have low idiosyncratic valuations but because their common-value element takes on an unexpectedly low value that falls short of the court's allocative price.

163. Although, as we note above, *supra* text accompanying note 151, Kaplow and Shavell fail to compare the property rule to the most efficient single-chooser rule because their example uses the less efficient chooser and nonoptimal damages.

164. Ayres & Talley, *supra* note 9, at 214; Epstein, *supra* note 9, at 2195; Rose, *supra* note 9, at 2180.

panoply of liability-rule implementations than merely the call-option rules of Calabresi and Melamed.

To get a fuller sense of why other implementations of the four basic allocations might be useful, it is useful in closing to speculate on how the foregoing analysis might be applied to breach of contract.[165] Contractual entitlements are normally protected by a single type of liability rule. The promisor has a call option on the promisee's entitlement to performance. In other words, through breach, the promisor can take the promisee's entitlement to performance and merely pay compensatory damages set at the court's best estimate of the promisee's expected value of performance. Thus, unless the court feels the situation warrants the extraordinary damage award of specific performance, the standard remedy implements a kind of Rule 2. Recharacterized in this way, we can see that contract law lags behind nuisance law — because while at least one court (in *Spur*)[166] has applied Calabresi and Melamed's Rule 4, contract courts do not typically allow promisees to choose between specific performance or damages.

Accordingly, contract remedies inefficiently place allocative authority solely in the hands of the plaintiff promisor. True expectation damages, as an implementation of a defendant-choice allocation, can induce defendants to breach efficiently, given their private information about the costs of breach. But the foregoing analysis suggests that, at times, the plaintiff will be better placed to decide whether performance should take place, or at still other times both the plaintiff and the defendant will make better allocative decisions (than either one or the other could do by themselves). The traditional expectation damage rule does a good job of harnessing the defendant-promisor's private information but does not harness the private information of the plaintiff-promisee. At times (particularly when the litigants have similar variances but different means), the court would do better to establish a dual-chooser allocation, announce an allocative price, and then give each party the opportunity to veto the defendant's breach.

Or, as an alternative, courts could give the promisees an option to increase the damages they would receive for the promisors' breach by making a firm offer to pay more for performance.[167] Thus, after a seller repudiated a contract through anticipatory breach, a court might give the buyer the option of offering the seller an additional amount to per-

---

165. To begin, the absence of contract law from Calabresi and Melamed's view of the cathedral is striking. The subject might have been omitted because contracts didn't fit their theory well. Calabresi and Melamed have been understood by subsequent scholars to have argued that liability rules tend to dominate when transaction costs are high but that property rules dominate when transaction costs are low. Contractual entitlements don't fit because they are protected by liability rules even though transaction costs were sufficiently low for the parties to enter into an initial contract.

166. Spur Indus., Inc. v. Del E. Webb Dev. Co., 494 P.2d 700, 708 (Ariz. 1972).

167. Ayres & Balkin, *supra* note 9, at 746.

form. The seller could either accept this amount and perform, or breach the contract and pay ordinary damages *plus the additional amount that the buyer offered to pay*.[168] Judges could effect this scheme by instructing the jury to determine ordinary expectation damages (without knowledge of the attempted modification) and then simply adding this augmentation to the jury's award. This mechanism would allow the plaintiff-buyers to signal credibly when their valuation of performance exceeded the provable mean valuation.[169] The defendant-sellers would then face a more authentic allocative price in deciding whether to breach.

Both contract and nuisance law provide rich settings in which to apply our analysis, but others exist as well. For example, many countries use "compulsory licensing" schemes to mitigate the bilateral monopoly problems of "blocking patents."[170] Compulsory licenses, of course, amount to a call-option implementation of a single-choice allocation. By now, however, grants of the allocative choice to the other side or to both sides in a blocking patent dispute obviously might yield superior results, and courts might then usefully consider pay-or-be-paid implementations to adjust the expected distribution. At the end of the day, bilateral disputes over rivalrous entitlements do not come close to describing the "cathedral" of allocative dilemmas facing courts, but they represent a substantial enough proportion of the world to justify our inquiry.

---

168. Courts should limit the buyer's ability to augment the seller's potential damages to circumstances in which performance can feasibly occur. Otherwise, buyers might make an inflated offer for performance (that is, an offer that exceeds their private valuation of performance) if they know that changed circumstances have rendered the seller's performance impossible. *See id.* at 747 n.137.

169. Buyers deciding how much extra to offer for performance would have a powerful incentive not to offer too much. If the offer grows too generous, the seller would simply accept it and perform.

170. *See* Gianna Julian-Arnold, *International Compulsory Licensing: The Rationales and the Reality*, 33 IDEA 349 (1987) (noting that blocking patents are one of the three most common conditions for compulsory licensing abroad). Countries that have compulsory licenses for blocking patents include Australia, China, France, Japan, the Netherlands, New Zealand and Switzerland. *See* ROBERT MERGES, INTELLECTUAL PROPERTY RIGHTS AND BARGAINING BREAKDOWN: THE CASE OF IMPROVEMENT INVENTIONS AND BLOCKING PATENTS 30-31 (1994). U.S. patent law also mandates compulsory licenses with regard to nuclear power and environmental engineering technologies. *See* Ayres & Talley, *supra* note 9, at 1093.