A PRIMER ON THE ECONOMICS OF CONFLICTS OF INTEREST

PETER SIEGELMAN*

ABSTRACT

There is a well-known conflict of interest between liability insurers and policyholders with respect to the decision to settle or litigate a claim. This short note provides a simple graphical explanation for the problem and grounds it in the way the structure of the parties' payouts drives their attitudes towards risk. An optional appendix links the insights to the elementary mechanics of financial options.

TABLE OF CONTENTS

* Phillip I. Blumberg Professor, University of Connecticut School of Law, peter.siegelman@uconn.edu. Thanks to Tom Baker, Bill Goddard, and, especially, Travis Pantin for careful reading and very helpful comments. I'm also grateful to the editors at the CILJ for their improvements to the piece. All errors and omissions are mine.

I. INTRODUCTION

This essay offers a graphical explanation of the conflict between insurers and policyholders in the decision about whether to settle or litigate a claim under a liability insurance policy with limits on coverage.[1] This complex and important topic has been the subject of voluminous litigation[2] and has attracted considerable attention from sophisticated legal scholars and policymakers.[3] My goal here is decidedly not to break any new ground. Rather, I offer a simple, visual, and hopefully, intuitive way of understanding *why* conflicts between insurers and policyholders arise and how they interact with other aspects of litigation such as the costs of suit and the "quality" or merits of the case. To that end, I've suppressed many of the institutional details that make the subject so tricky.

II. LITIGATION AND LIABILITY INSURANCE

A. Basic Mechanics

Liability insurance protects policyholders from the risk that they will have to pay (e.g., in tort) for harm they have imposed on a third party. Such policies cover the insured defendant for the amount that a court awards the victim; if the parties settle the lawsuit instead of litigating it to a final judgment, the insurer is also responsible for the settlement amount. Typically, the policy also covers the cost of defending against a lawsuit brought against the policyholder.

---

[1] As a bonus, the analysis maps neatly into the elementary theory of financial options. For readers familiar with basic option theory, the figures presented below will be easily recognizable; for those who are not, the economic insights will be derived independently. An appendix summarizes the relevant option analysis.

[2] *See, e.g.*, Crisci v. Sec. Ins. Co. of New Haven, 426 P.2d 173 (Cal. 1967). As of July 25, 2023, *Crisci* has been cited in 517 state and federal judicial opinions and 330 law review articles.

[3] RESTATEMENT OF L. OF LIAB. INS. § 24 (AM. L. INST. 2019). The scholarly literature includes, but is not limited to, Robert E. Keeton, *Liability Insurance and Responsibility for Settlement*, 67 HARV. L. REV. 1136 (1954); Kent D. Syverud, *The Duty to Settle*, 76 VA. L. REV. 1113 (1990); Alan O. Sykes, *Judicial Limitations on the Discretion of Liability Insurers to Settle or Litigate: An Economic Critique*, 72 TEX. L. REV. 1345 (1994); Tom Baker, *Liability Insurance Conflicts and Defense Lawyers: From Triangles to Tetrahedrons*, 4 CONN. INS. L. J. 101 (1998); Ezra Friedman, *The Value of a Statistical Judgment: A New Approach to the Insurer's Duty to Settle*, NW. L. & ECON. SERIES, no. 15–03, Dec. 2014, at 1, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2553439.

Crucially, however, the policies almost always contain limits on coverage. These are dollar amounts beyond which the insurer is not responsible. This means that if the award (or settlement) is greater than the policy limit, any excess must be paid by the policyholder. Because the insurer is always responsible for at least some payouts, the insurer usually[4] has control over most aspects of the litigation strategy, including the decision about whether to settle a claim against the policyholder and how much to spend in defending it. For now, we will unrealistically assume that litigation is costless, but we'll relax that assumption later.

Suppose that a policyholder (PH) has liability insurance with a coverage limit of $120,000, with no deductible.[5] He is sued by a plaintiff (P) for a covered slip and fall injury that occurred on his property.[6] Suppose the lawsuit is definitely going to trial and consider the insurer's payout as a function of the amount awarded by the jury. For any amount less than the $120,000 policy limit, the insurer pays 100%, so each dollar of award below this limit means a full dollar of loss for the insurer. If the award is greater than the limit, the policyholder is responsible for the excess, so the worst that can happen from the insurer's perspective is that it pays the policy limit. Graphically, this is depicted by the red line in Figure 1, which shows the insurer's payout increasing[7] dollar for dollar for judgments below the limit but then flattening out at the $120,000 limit, which is precisely what a limit is designed to do.[8]

The policyholder's payout (as a function of the jury award or settlement amount) is graphed by the blue line in Figure 1. Low awards or settlements—anything below the coverage limit—cost him nothing since they are entirely paid for by his insurer. He is only responsible for any amount above the $120,000 limit. Therefore, his payout is horizontal (at zero) for awards below the limit, and then increases dollar-for-dollar with

---

[4] Though not always. *See* Douglas R. Richmond, *Liability Insurance and the Duty to Pay Defense Expenses Versus the Duty to Defend*, 52 TORT TRIAL & INS. PRAC. L. J. 1, 8 (2016); Charles Silver, *Basic Economics of the Defense of Covered Claims*, *in* RESEARCH HANDBOOK ON THE ECONOMICS OF INSURANCE LAW 438, 438 (Daniel Schwarcz & Peter Siegelman eds., 2015).

[5] *See infra* Section III.B for analysis of coverage with a deductible.

[6] Purely to assist in keeping track of the parties, I'll refer to the policyholder as "he" and the plaintiff as "she." The insurer is an "it."

[7] Since we are dealing with payouts or losses by the insurer or policyholder, the amounts are all negative. It is the (absolute) *magnitude* of the insurer's payout that increases as the amount awarded rises (below the policy limit). Those minus signs can be tricky.

[8] The phrase "dollar-for-dollar" means that the slanted part of the insurer's payout function, below the limit, has a slope of -45°.

any amount that remains.[9]

Finally, the total amount obtained by the plaintiff is just the combined amount paid by the insurer and the policyholder, which is the vertical sum of the two parties' payouts. This forms the straight black line (with slope (-45°)) in Figure 1. That makes sense, since every dollar the plaintiff receives must come from one or the other of the two possible sources of payment. In our simple example, the plaintiff shouldn't notice or care that the first $120,000 of payout comes from one source and anything above that from another.[10]

[9] That is, an award of $120,001 costs the policyholder $1.

[10] This simplifying assumption should not be understood as denying the existence of differences between "insurer money" and "policyholder money." In the real world, such differences can be important. For instance, defendants may not have the wealth to pay awards in excess of the policy limits, while insurers generally do. *See, e.g.*, Steven Shavell, *The Judgment Proof Problem*, 6 INT'L. REV. OF L. AND ECON. 45 (1986). And there can be *moral* differences between insurer money and defendant money. *See,* Tom Baker, *Blood Money, New Money, and the Moral Economy of Tort Law in Action*, 35 L. & SOC. REV. 275, 281, 301 (2001) (demonstrating that plaintiff-side tort lawyers often believe it is illegitimate to seek damages in excess of the policy limits, and refer to money paid out of defendant/policyholder pockets as "blood money."). Thanks to Travis Pantin for this insight.

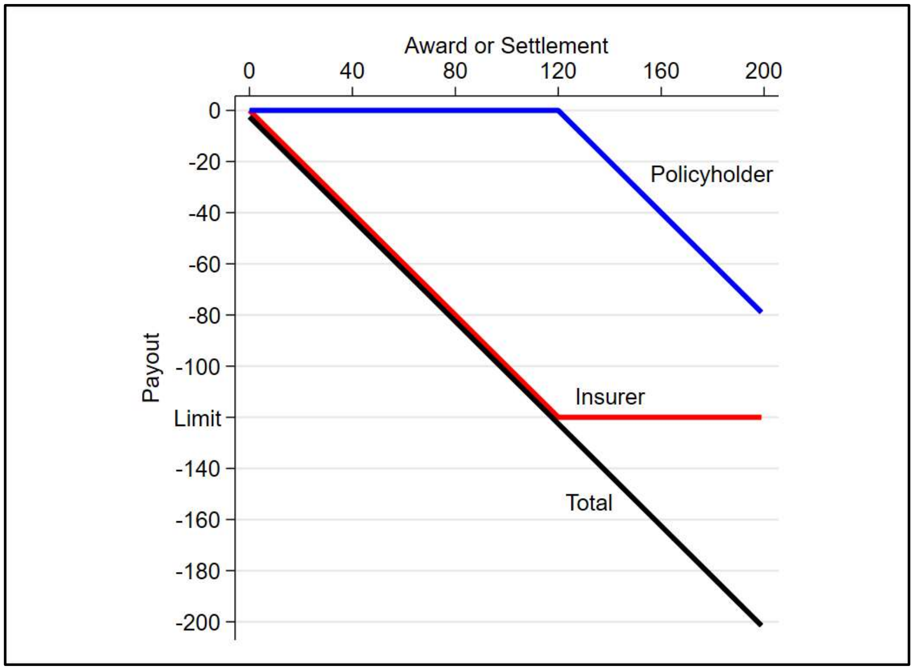

Figure 1: Insurer, Policyholder and Total Payouts as Functions of Award or Settlement (Policy Limit = $120,000)

B. ANALYSIS

Hidden in Figure 1 are some deep insights about the differences between the insurer's and the policyholder's feelings about risk, which drive the well-known conflict between the parties that arises in situations where the award may exceed the policy limit. To uncover these principles, we need to introduce risk into our analysis. Instead of assuming that the outcome in the underlying litigation is known for sure, let's suppose that it is uncertain. More specifically, suppose everyone agrees there is a 50% chance that the plaintiff will prevail at trial and be awarded $200,000 and a 50% chance that she will lose and receive nothing. This dispersion of possible outcomes is what we mean by risk.

Remember, we're still assuming there are no trial costs, so all that is at stake is the amount of the award. Finally, let's also assume for the time being that the insurer has complete control over the litigation, meaning that it alone decides whether to accept any settlement offer or risk going to trial instead.

1. The Plaintiff's Perspective

To understand how the parties view the litigation when the outcome is uncertain, it will be useful to start with the "expected value" of the lawsuit to the plaintiff (P).[11] Expected value (EV) is the "average" outcome, where average is understood as the "probability-weighted average of the possible outcomes." That is, letting the *P* subscript denote the plaintiff:

$$EV_P = Prob.(lose) \times Amt. Rec'vd\ if\ Lose + Prob.(win) \times Amt. Rec'vd\ if\ Win \quad (1)$$

$$= \frac{1}{2}(\$0) + \frac{1}{2}(\$200{,}000) = \$100{,}000.^{12} \quad (1)$$

Purely for convenience, we'll assume that the plaintiff is risk-neutral. That is, she is indifferent between getting the expected value of a gamble or the actual gamble itself.[13] Here, because the expected value of the lawsuit to the plaintiff is $100,000, she would require at least that much to settle the case and avoid trial. Anything less and she would prefer to roll the dice at trial.

2. The Insurer's Perspective

Given the uncertain outcome, how will the insurer feel about taking this case to trial? Remember, the insurer's maximum exposure is capped at the policy limit, so when it computes the expected value of going to trial, it knows it will not have to pay the full amount of the award if the plaintiff wins; at worst, it will only be liable for the policy limit. Because it will never have to pay more than the policy limit, the *insurer's* expected value of going to trial is:

$$EV_I = -[\tfrac{1}{2}(\$0) + \tfrac{1}{2}(\$120{,}000)] = -\$60{,}000. \quad (2)$$

---

[11] The expected value of the lawsuit to the plaintiff's combined opponents—the insurer and policyholder *taken together*—is simply the negative of this amount.

[12] Note that since the plaintiff will either win or lose, the total probability of those outcomes must sum to 1.

[13] Someone who is risk-neutral would be indifferent between: (a) a gamble that wins $1 if a (fair) coin comes up heads and loses $1 if the coin comes up tails; and (b) $0, for sure. They would take the bet if you offered them 1 cent; they'd refuse to pay 1 cent to play.

Any judgment (or settlement) in excess of that limit lies on the *flat* portion of the insurer's payout diagram in Figure 1 (in red), which means that from the insurer's perspective, a judgment of $200,000 is no worse than a judgment of $120,000—it will owe the same $120,000 in either case. Given this, the insurer will reject any settlement offer from the plaintiff that costs it more than $60,000, since a payment of anything more than that would be worse (on average) for the insurer than going to trial.[14]

Graphically, the expected payout by the insurer of a lawsuit that has a 50/50 chance of awarding $0 or $200,000 (but is subject to a $120,000 policy limit) is represented in Figure 2. It is just the midpoint of the dashed line connecting the points (0,0) and (200,-120), denoted by the X. Notice that the insurer's *expected* payout when facing this gamble ($60,000) is smaller in absolute magnitude than (lies *above*) the actual payout that would arise from a settlement or judgment of $100,000. Since that amount is less than the policy limit, the insurer would have to pay all of it. By taking a gamble on trial, the insurer is likely to leave itself at least as well off as if it had to pay the full $100,000 expected value in settlement.

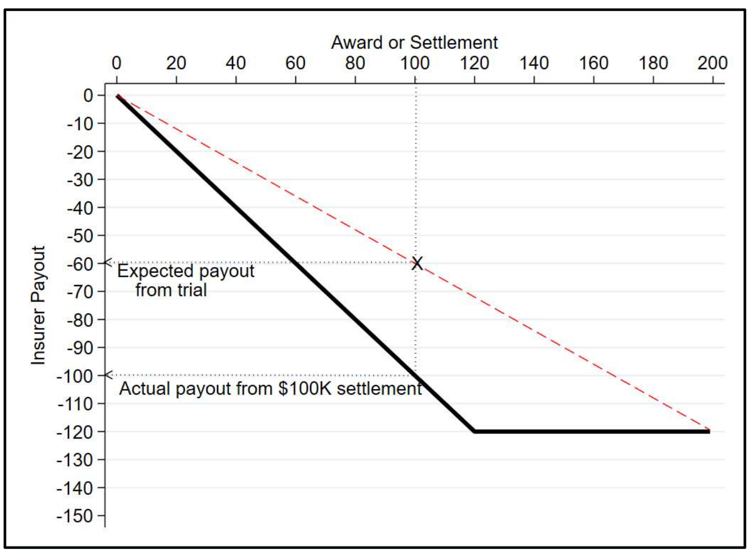

Figure 2: Insurer's Payout and Expected Value of Lawsuit
with Liability of 0 or 200, each with 50% Probability
(policy limit is 120)

[14] Of course, the insurer might have all kinds of additional reasons to prefer litigating, including establishing a reputation as a tough negotiator in future cases. But we abstract from those motives to simplify the analysis.

It should be clear, then, that when the insurer controls the decision, there is no possibility of settling this case. The plaintiff will reject any offer to settle that leaves her worse off than going to trial would (on average), which cashes out to a \$100,000 demand. The insurer will refuse to pay anything more than what it would expect to pay if the case went to trial (\$60,000), so there is no bargain to be struck here—litigation is unavoidable.

Crucially, this is not simply a feature of the particular numbers chosen. Rather, it is dictated by the shape of the insurer's payout function: the fact that it flattens-out at the policy limit that caps its exposure means that whatever part of the award is above the policy limit "doesn't count" from the insurer's perspective. Put differently, the insurer will always gain by taking a chance on litigation rather than choosing the sure thing—settling (at the claim's expected value). That's because there is a structural asymmetry in the way that awards to the plaintiff translate into payouts by the insurer. An extra dollar of any "large" award is free to the insurer, because once the award surpasses the policy limit, every marginal dollar is paid by the policyholder. The cost of a "small" award, by contrast, is entirely borne by the insurer, so an additional dollar awarded to the plaintiff costs the insurer a full extra dollar.

To see how this plays out, consider an alternative lawsuit that will result in an award of either \$1,000 or \$199,000, each with 50% probability. It should be clear that this second suit has the same expected *total award* as the first one: ½(\$1,000)+½(\$199,000) = \$500 + \$99,500 = \$100,000. Although it has the same expected award, the second suit is less risky than the first. In finance, risk is typically measured by the standard deviation (SD), which captures the dispersion of outcomes around their average. For the first lawsuit, the SD of the award is 100, while for the second it is 99.[15] Comparing standard deviations reveals what should be clear intuitively—the second lawsuit is less risky than the first, because its outcomes are clustered more closely around the average or expected value.

Notice that the insurer will always prefer to face the first suit than the second because *its* own expected payout is larger under the second—\$60,500 (= ½(\$1,000) + ½(\$120,000)), versus \$60,000 for the first. To be sure, the second suit does have lower maximum exposure (\$199K vs \$200K). However, that "savings" does the insurer no good since it comes from over-the-limit dollars the insurer was never going to pay in the first place. On the flip side, the second suit has \$1,000 of guaranteed exposure (win or lose), while the first suit costs nothing if the insurer wins. All of this implies that the insurer's payout is structured so that its expected payment actually falls

[15] $SD_1 = \sqrt{(200{,}000 - 100{,}000)^2+(0 - 100{,}000)^2} = 100$. $SD_2 = \sqrt{(199{,}000-100{,}000)^2+(1{,}000-100{,}000)^2} = 99$.

when the risk rises. Naturally, the insurer *likes* that increase in risk.[16]

3. The Policyholder's Perspective

We've assumed thus far that it is the insurer who gets to decide whether to take the safe option and settle or risk going to trial. But we might imagine that this right belongs instead to the policyholder. How will he evaluate the situation?

Remember, the plaintiff has a 50% chance of winning at trial, and if she wins, she'll be awarded $200,000. She should thus be willing to accept anything more than the case's expected value of $100,000 to settle it. But given that the policyholder is only responsible for payouts above the policy limit, his expected value of going to trial is:

$$EV_{PH} = -[\tfrac{1}{2}(\$0) + \tfrac{1}{2}(\$200{,}000 - \$120{,}000)] = -\$40{,}000. \quad (3)$$

This means that the policyholder would be willing to spend up to $40,000 of his own money to avoid going to trial, since that's his expected loss. Consider a settlement offer by the plaintiff for $100,000. That's well below the policy limit, so if the policyholder controls the decision, he would happily agree to that amount. It leaves him with no out-of-pocket costs at all, since the entire settlement will be covered by his insurer.[17]

Figure 3 provides a graphical intuition for this result. Awards below the policy limit lie on the flat part of the policyholder's payout function. For

---

[16] Finance geeks might appreciate that the insurer's payout is a *convex* function of the amount awarded. A straightforward way to define a convex function is that it is one whose value at the midpoint of every interval in its domain is less than or equal to the average of its values at the ends of that interval. For the function

$f(\,)$, which maps awards to insurer payouts, to be convex it must be true that for any awards a and b:

$f(\tfrac{1}{2}a + \tfrac{1}{2}b) \leq \tfrac{1}{2}f(a) + \tfrac{1}{2}f(b)$.

This is precisely the case with the insurer's payout function. The key finance insight is that payout convexity is associated with a preference for risk (i.e., the opposite of risk-aversion). Someone with a convex payout structure always prefers a gamble to its expected value: that is, they always want to take a bet on a coin flip that wins $1 on heads and loses $1 on tails. Note that the insurer's risk-preference is not generic—it applies only to this particular problem and arises only from the structure of the insurer's payout function. It could well be that the insurer is *generally* risk-neutral, even though *in this context* it will prefer to take a gamble on litigation rather than settle for the expected award in that litigation.

[17] Indeed, the policyholder would willingly agree to any total settlement of less than $160,000 (with $120,000 contributed by the insurer) rather than face trial.

him, an award of $100,000 with certainty is costless; but going to trial has an expected payout (denoted by the X) of $40,000 (midway between the two possible payouts of $0 and $80,000.

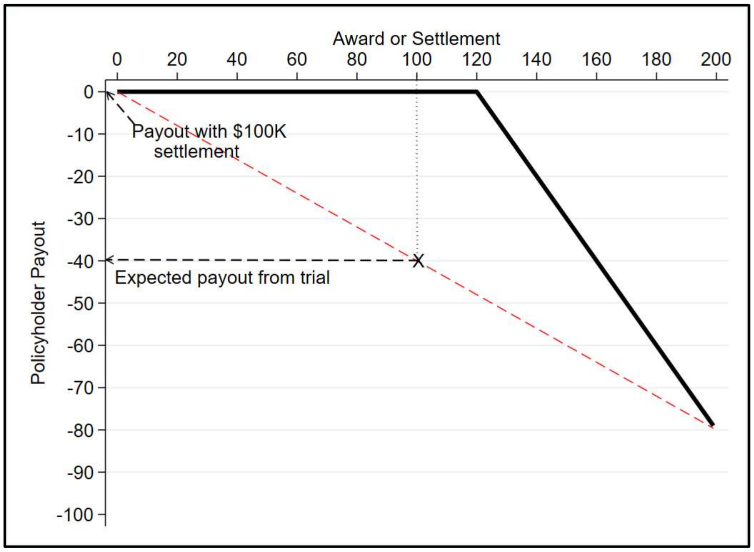

Figure 3: Policyholder's Payout and Expected Value of Lawsuit with Liability of 0 or 200, each with 50% Probability (policy limit is 120)

Just as with the insurer, the policyholder's preferences for risk are inherent in the structure of the payout function, though they take the opposite form. To see why, consider again the alternative lawsuit that generates liability of either $1,000 or $199,000, each with 50% probability. The plaintiff's expected award from this suit is still $100,000. But the policyholder will prefer to face the second lawsuit rather than the original one. In the second suit, the policyholder's expected payout is (½($0) + ½($199,000 - $120,000)) = $39,500, $500 less than before. The worst outcome from the policyholder's perspective is not as bad ($79,000 of over-the-limit exposure, vs $80,000), while the best outcome is unchanged, since the additional $1,000 minimum payout will always be paid entirely by the insurer. The second lawsuit is less risky (has a smaller SD), and the policyholder prefers it to the first one because he is risk averse.[18]

---

[18] That is, the policyholder's payout is a *concave* function of the award or settlement. This reverses the key finance insight from the previous note: payout concavity is associated with risk-aversion, rather than risk-preference. Someone with a concave payout function would always refuse to bet on a coin flip that wins $1 on heads and loses $1 on tails. Note that, again, we are not making a "global" statement about the policyholder's risk preferences. The policyholder's underlying

The basis of the conflict between insurers and policyholders can therefore be seen as arising from a difference in attitudes towards risk that are caused by each side's payout function. The insurer's payout function is capped at the $120,000 policy limit, so it prefers riskier outcomes; the policyholder's payout function is costless until the limit is hit, so he prefers less-risky outcomes. Settlement is less risky than trial, and the policyholder and insurer have opposite views of whether this reduction in risk is good or bad.

4. Symmetry?

It might be tempting to think that the insurer's behavior described earlier—turning down the plaintiff's settlement offer at the case's expected value—is "illegitimate," since it entails gambling with the policyholder's money. If the plaintiff loses, nobody pays anything; if the plaintiff wins, the insurer pays the policy limit, and the policyholder is stuck with the remaining $80,000. That characterization is logically correct. However, the same logic applies in reverse when the policyholder controls the settle/litigate decision and agrees to settle the case for its expected value (or more). By settling for an amount *under* the policy limit,[19] the policyholder would be spending the insurer's money.

The real issue is that when the responsibility for compensating the plaintiff is split between the two parties, but the decision about settlement is allocated (exclusively) to one of them, the structure of the problem guarantees that this party will end up "playing with the other's money." That is, for any given award (or settlement) in Figure 1, one party will always be operating on the flat part of its payout function, where additional amounts come out of someone else's pocket and cost that party nothing.

At this point, however, a caution is in order. Just because the *math* is in some sense symmetric doesn't mean that the *parties* should be treated symmetrically. That is a normative conclusion, about which the analysis is silent. In particular, a preference for one party over another in this situation might well take into account the parties' relative sophistication, risk-tolerance, and ability to spread risk.[20] None of those elements are present in

---

utility function (which maps her wealth into her utility) could make her risk-loving in most contexts. But in this narrow context, the payout structure means that the policyholder will always prefer less risk to more.

[19] More precisely, since the policyholder would be willing to spend up to $40,000 of his own money to avoid going to trial, any total settlement less than $160,000 (of which $120,000 is contributed by the insurer) would be preferable to going to trial.

[20] Communication with Tom Baker, William Maul Measey Professor of L., U. Pa. Carey L. Sch., who has stressed that insured defendants are almost *by definition* generically risk-averse, almost always undiversified, and typically exert no control

the simple model sketched above, which therefore sheds no light on how the conflict should best be handled. If it does anything, it only illuminates the structure of the conflict.[21]

C. THE IMPORTANCE OF REASONABLENESS

One thing the model can shed some light on is the importance of reasonableness in assessing the insurer's settle/litigate decision. As noted earlier, that decision typically belongs to the insurer.[22]

Consider a different variation of the original facts above. The plaintiff's lawsuit still pays $200,000 if she prevails, and pays nothing if she loses. She again makes an offer to settle the lawsuit for $100,000. But now suppose that its probability of success is only 20% rather than 50%. That means the plaintiff's expected value from trial is now:

$$EV_P = 0.2(\$200{,}000) + 0.8(\$0) = \$40{,}000. \tag{4}$$

The insurer's expected payout from trial becomes:

---

over policy language. So various "contract interpretation" risks (including the risk of a conflict with one's insurer that is not covered by the policy language) might best be assigned to the insurer. When that occurs, competition should increase the premium paid for coverage, but given the policyholders' risk aversion, the tradeoff (higher premium for more coverage) will typically be welfare-enhancing.

[21] One intriguing proposal—which, however, seems not to have gotten much traction—comes from Richard Squire, who offers a relatively simple structural solution to the misaligned incentives: let each party separately resolve its slice of potential liability with the plaintiff. Richard Squire, *The Artificial Collective-Action Problem in Lawsuits Against Insured Defendants, in* RESEARCH HANDBOOK ON THE ECONOMICS OF INSURANCE LAW, *supra* note 4, at 461. Under his proposal, the policyholder would be free to settle his potential individual liability to the plaintiff, but this would not impact the insurer's ability to proceed to trial to determine what percentage of its policy limits it owes to the plaintiff. Squire's approach eliminates the capacity of either the insurer or policyholder to shift exposure to liability onto the other, but it does have some downsides. One is that it increases risk to policyholders, who might be asked to contribute some amount to settlement more often than under the "ignore the limits" rule favored by the Restatement of Liability Insurance. RESTATEMENT OF THE L. OF LIAB. INS. § 24 rep. note b (AM. L. INST. 2019). Savvy policyholders would recognize this ex ante and demand lower premiums for a Squire-rule policy, but of course many policyholders are not savvy. An additional wrinkle is that the "ignore the limits rule" might in some cases work to the advantage of plaintiffs. Sykes, *supra* note 3. That should tend to make a Squire-rule policy, which avoids this problem, somewhat cheaper in equilibrium.

[22] *See* discussion *supra* Section II.B.

$$EV_I = -(0.8(\$0) + 0.2(\$120{,}000)) = -\$24{,}000. \quad (5)$$

The insurer will still prefer to go to trial rather than settle this case for the \$100,000 the plaintiff has demanded. But here, the plaintiff's \$100,000 ask grossly (2.5✕) exceeds the true value of the litigation, and the insurer's decision not to settle seems entirely appropriate.

Law and policy will therefore need to do more than simply require the insurer to settle whenever there is risk of an award that is above the policy limit. That rule would be too crude, because it fails to capture the difference between this example and the previous one—sometimes, plaintiffs demand much more than they could expect to win at trial, and it would make no sense for the law to require the insurer to acquiesce in that situation.

Given all this, there are only three possible rules that the law could embody. First, it might impose no constraints at all on the insurer's settlement decision, leaving it up to the insurance contract, as "negotiated" by the parties themselves.[23] The analysis above demonstrates that this will inevitably generate some cases where the insurer ends up litigating with the policyholder's money and exposing its insured to serious liability. The famous *Crisci* case[24] is an example of what can go wrong here. Still, there might be a case for no regulation if one believed that policyholders knew about the potential conflict and could negotiate for a lower price that reflected the higher risk they face.

A second choice might be to simply require settlement; but that is clearly unattractive for reasons highlighted above. A final alternative would be to impose a "soft" requirement of reasonable settlement behavior. This is precisely the approach adopted by The Restatement of the Law of Liability Insurance, which takes the position that "[w]hen an insurer has the authority to settle a legal action brought against the insured . . . and there is a potential for a judgment in excess of the applicable policy limit, the insurer

---

[23] "Negotiated" is in quotation marks because in many contexts, it makes no sense to suppose that the parties actually bargain over any of the terms in an insurance contract. *See* RESTATEMENT OF THE L. OF LIAB. INS. § 2, cmt. d (AM. L. INST. 2019) (explaining that insurance contracts are standard forms, meaning that policyholders can choose coverage only by selecting from forms provided by the insurer. "Even in the commercial insurance market, the vast majority of insurance policies are standard-form contracts."); *see also* RESTATEMENT (SECOND) OF CONTRACTS § 211, cmt. d (AM. L. INST. 1981) (concluding that "[a] party who makes regular use of a standardized form of agreement does not ordinarily expect his customers to understand or even to read the standard terms."). Thanks to James Hallinan for these references and suggestions.

[24] Crisci v. Sec. Ins. Co. of New Haven, 426 P.2d 173 (Cal. 1967).

has a duty to the insured to make reasonable settlement decisions."[25] Of course, this rule may raise difficult factual questions about whether the case would be worth litigating if the insurer bore the risk of an adverse outcome, but that is always the case whenever a reasonableness standard is adopted.

The simple model presented here does not illuminate exactly how the duty to make reasonable settlement decisions should be characterized.[26] But it does suggest that if courts are not going to adopt a completely laissez-faire regime with respect to these conflicts, it will be difficult to do better than some version of a reasonableness standard: qualitative dimensions (such as how strong was the plaintiff's case) are necessarily at play, and it will be hard to formulate a crisp rule that covers all these dimensions.

III. FURTHER TWEAKS

A. LITIGATION COSTS

So far, we have assumed that litigation and settlement are costless to all parties. That simplification made sense as a way to focus on the essential structure of the problem, but it is obviously wrong. Indeed, avoiding the cost of trial is presumably a key motive for the parties to settle: doing so minimizes payments to others (e.g., lawyers) which the parties can keep for themselves if they can negotiate a settlement.

Does recognizing that trials are costly change the analysis above? Unfortunately, there is no longer a simple graphical analysis, but the answer is, "maybe" (if trial costs are sufficiently high).[27] To see why, suppose that

---

[25] RESTATEMENT OF THE L. OF LIAB. INS. § 24(1) (AM. L. INST. 2019). The Restatement adopts the view that the "disregard the limits" rule is the appropriate standard for reasonable behavior. That rule "has . . . become the most common test for determining whether an insurer gave 'equal consideration' to its insured's interests in duty-to-settle cases," and requires that the insurer accept any settlement offer that "a prudent insurer without policy limits would have accepted . . . ." *Id.* at rep. note b (citations omitted). Note that the Restatement's "disregard the limits" is a default rule: the parties can contract for something else if they choose to do so, and the Restatement's formulation only operates when the parties are silent.

[26] It also ignores many other aspects of the problem. For example, policy limits are not chosen at random, and a rule that said that insurers were free to settle or litigate as they choose would presumably put pressure on policyholders to select higher limits for fear of a *Crisci* situation arising. 426 P.2d at 177–78. A full analysis is vastly more complicated than the simple story sketched here.

[27] The reason there is no longer a simple graphical analysis is that the introduction of litigation costs breaks the identity between what the plaintiff receives

as above, the policyholder has coverage with a limit of \$120,000.[28] The plaintiff's claim still has a 50% chance of a \$200,000 award at trial and a 50% chance of \$0. But we now assume that it costs \$50,000 for each side to litigate the lawsuit. Settlement, however, is costless. The presence of litigation costs means that what one side pays is no longer identical to what the other side receives. That identity still holds *if the case settles* (and litigation costs are avoided); but if the case goes to trial, there is now a \$100,000 "wedge" between what's paid and what's received, with the gap accounted for by each side's legal fees. Let's consider what happens when the insurer controls the settle/litigate decision, which was the source of the conflict we described earlier.

Now, the insurer's expected payout from going to trial is:

$$EV_I = -[\tfrac{1}{2}(\$50{,}000) + \tfrac{1}{2}(\$120{,}000)] = -\$85{,}000. \quad (6)$$[29]

This is more costly than in the previous example, since there are litigation expenses incurred (for which the insurer is responsible) even when there is a pro-defendant verdict and no actual liability.

Note that the plaintiff faces litigation costs as well. Her expected value of litigation is now:

---

and what the insurer and policyholder together pay—we now must keep track of payments to a party.

[28] We will assume that defense costs are "within limits," meaning that all such costs count against the policy's overall coverage limit, cutting into the amount available to pay for any actual award. Thus, when the plaintiff loses, the insurer pays \$50,000 in litigation expenses, but nothing to cover any judgment. When the plaintiff wins, the insurer pays the entire policy limit (consisting of \$50,000 in litigation costs and \$70,000 towards the judgment). If the policy were written with defense costs treated separately (excluded from policy limits), the analysis looks essentially the same. The insurer's expected trial cost is:

$EV_{I'} = -[\frac{1}{2} \times \$0 + \frac{1}{2} \times \$120{,}000] - \$50{,}000 = -\$110{,}000$.

That is, the insurer always pays litigation costs in addition to the full policy limit of \$120,000. So the insurer has even more exposure than before and an even stronger reason to settle the case. *See* discussion *infra* Section III.A.

[29] We focus on the insurer here. But the policyholder's expected value from going to trial is now:

$EV_{PH} = -[\frac{1}{2}(\$0) + \frac{1}{2}(\$200{,}000 - (\$120{,}000 - \$50{,}000))] = -\$75{,}000$.

The \$120,000 limit is effectively lowered by the \$50,000 in litigation costs that are incurred if the case goes to trial, so there is only \$70,000 available to pay the plaintiff, with the policyholder responsible for the remaining \$130,000 of the award.

$$EV_P = [\tfrac{1}{2}(-\$50{,}000) + \tfrac{1}{2}(\$200{,}000 - \$50{,}000)] = \$50{,}000.^{30} \quad (7)$$

The plaintiff would settle for anything more than $50,000, while the defendant would settle for anything less than $85,000. There is now a bargaining surplus available to the parties if they can reach a settlement, so rational litigants will want to settle the case, even when the insurer controls whether to go to trial.

This means that the policyholder will—almost accidentally—receive some "protection" from the insurer's "excessive" willingness to risk going to trial, merely because settlement offers a reason to avoid incurring those expenses. That's true even for an insurer who faces a policy limit that would otherwise lead it to prefer trial to settlement. Of course, this will not always be the case: if trial costs are only $10,000, for example, the insurer will expect to pay $65,000 at trial, the policyholder will expect to receive $90,000, and there is no mutually beneficial deal to be struck that avoids litigation. The insurer/policyholder conflict will not be eliminated under these circumstances.

B. Deductible

So far, we have analyzed the problem of shared payouts on the assumption that the insurer is responsible for all the payout up to the policy limit, while the policyholder is only responsible for those payouts above the limit. But insurance policies frequently contain a deductible, which is just a requirement that the policyholder is responsible for the "first" dollars of any payout up to the deductible amount; after that, the insurer pays any part of the award or settlement until the limit is reached, at which point the policyholder is again responsible for all payouts.[31]

Figure 4 illustrates how the presence of a deductible (here, assumed to be $20,000) changes the analysis. (We maintain the assumption of a

---

[30] When the plaintiff loses at trial, she incurs $50,000 in expenses and receives nothing. When she wins, she incurs $50,000 in expenses and receives a judgment of $200,000, for a net of $150,000.

[31] Like so much else in the economics of insurance, the theory of optimal deductibles was first explored by Kenneth Arrow. *See* Kenneth J. Arrow, *Uncertainty and the Welfare Economics of Medical Care*, 53 AMER. ECON. REV. 941, 960 (1973) (concluding that when the insurer charges "a fixed-percentage loading above the actuarial value for its premium[,] . . . the most preferred policy from the point of view of an individual is a coverage with a deductible amount; that is, the insurance policy provides 100 per cent coverage for all . . . costs in excess of some fixed-dollar limit."). Interestingly, Arrow's analysis implies that coverage limits are not optimal, at least in the relatively simple model he presents.

$120,000 coverage limit, but that limit now caps the insurer's payout $120,000 on *top* of the $20,000 paid by the policyholder.) Figure 4 shows that the insurer's payout (in red) and the policyholder's payout (in blue) now contain not one, but two, "kinks." The insurer pays nothing when the award or settlement is less than the $20,000 deductible, pays 100% of the next $120,000 (that is, up to $140,000 in award), and nothing thereafter. The policyholder's payouts are just the mirror image: awards of $0 to $20,000 are paid entirely out of pocket with no insurer contribution. After that $20,000 has been paid, any additional amounts up to $140,000 are solely the insurer's responsibility, so the policyholder pays nothing. But the part of any award greater than $140,000 is still paid entirely by the policyholder.

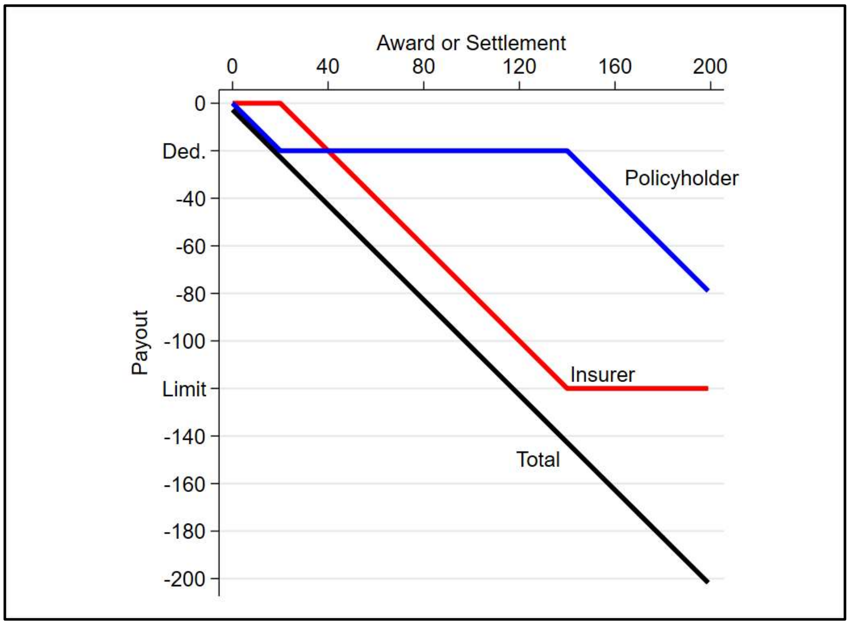

Figure 4: Insurer, Policyholder and Total Payout as Functions of Award or Settlement with Deductible of 20 (Policy Limit = 120)

The presence of a deductible changes the attitude of both parties towards some—but, interestingly, not all—risks. To see why, consider again a variation on our earlier lawsuit, in which the judgment at trial is either $200,000 or $0 (each with a 50% probability), where the policy limit is $120,000 as before. But now, suppose there is a deductible of $20,000. In this case, the presence of the deductible does *not* change the expected value of the payout to either party: The insurer expects to pay ½($0) + ½($140,000

- \$20,000) = \$60,000, as before;[32] the policyholder also expects to pay ½(\$0) + ½[(\$200,000 - \$140,000) + \$20,000] = \$40,000, as before. In this instance, the deductible alters nothing. When the plaintiff wins, the policyholder pays the "first" \$20,000 and the "last" \$60,000, instead of paying the "last" \$80,000 (as was the case without the deductible), but that is of no consequence.

However, the deductible does change attitudes towards risk, and preferences for trial or settlement. To see why, we need to consider the more general case. The insurer's expected payout in the presence of a deductible (D) and a policy limit (L) is given by:

$$EV_I = -\frac{1}{2}Max[V_{lo} - D, 0] - \frac{1}{2}Min[V_{hi} - D, L], \tag{8}$$

where V_{lo} and V_{hi} denote the smallest and largest verdicts, respectively. Similarly, the policyholder's expected payout in this situation is:

$$EV_{PH} = -\frac{1}{2}Min[V_{lo}, D] - \frac{1}{2}Max[V_{hi} - L - D, 0]. \tag{9}$$

Figure 4 illustrates the two "kinks" in each party's payout function, and these kinks generate the more complicated formulas above. The presence of the kinks means that each party's payout function changes shape—each is convex over some ranges and concave over others,[33] which in turn implies that each party's attitude towards risk now depends on the amount of the possible verdicts/settlements at issue. That is, each side is risk averse for some risks, risk-neutral for others, and risk-loving for yet others.

For example, consider a lawsuit against the policyholder that will either pay the plaintiff \$0 if the defendant prevails or \$30,000 if the plaintiff does, each with a 50% probability. Using equations (8) and (9) above, we can see that the insurer pays \$10,000 (the verdict in excess of the deductible) if the case goes to trial and the plaintiff wins. Since that happens 50% of the time, that's an expected value of -\$5,000. The policyholder pays \$20,000 (the entire deductible amount) if the plaintiff wins at trial, so his expected cost of trial is, similarly, one half that amount, or -\$10,000. (Of course, neither pays anything if the plaintiff loses.) The plaintiff's expected value from going to trial is \$15,000 (½(0) + ½(\$30,000)), so suppose she makes an offer to settle the case at that amount. That offer requires the insurer to pay nothing, since the total is less than the deductible. Conversely, the policyholder would have

---

[32] The \$140,000 reflects the fact that the insurer's payout is capped at the policy limit, but it doesn't start paying anything until the policyholder has paid the first \$20,000 in damages, using up her deductible.

[33] *See supra* notes 16, 18.

to pay the entire $15,000 settlement. But since that sum is more than his $10,000 expected cost from going to trial, he will want to reject the settlement offer and go to trial, while the insurer would obviously prefer to accept it.

Notice the role reversal here: For low-value claims in the presence of a deductible, the *policyholder* is the one who prefers the riskier alternative of trying the case while the insurer prefers the settlement of a claim in excess of the policy limit. While this scenario is theoretically possible, it seems unlikely to occur as a practical matter. If there are any costs of litigation, this kind of small-value claim is unlikely to be worth litigating in the first place. So this conflict is unlikely to arise in practice. Moreover, by definition, the stakes are small here, so the consequences of any conflict for the party who is not in control are not so severe as in the high-stakes example with which we began. And large risks—the kind we care most about—will still be generally subject to the same conflict of interest between insurer and policyholder as when there is no deductible, as shown earlier.

IV. CONCLUSION

In a lawsuit covered by liability insurance, policy limits that cap the insurer's exposure create well-known conflicts of interest between insured policyholders and their insurers. This short note explains why such conflicts are rooted in the structure of the problem, which shapes the parties' attitudes toward risk. The graphical analysis reveals when and why insurers prefer the risk of litigation and policyholders prefer to settle.

V. APPENDIX

This Appendix repackages the insights of the main text in terms of financial options. There is a clear and obvious parallel between the structure of financial options and the structure of the parties' payouts in the settle/litigate decision described in the main text.

A. OPTION BASICS

A financial option is simply the right—but not the obligation—to buy or sell an asset at a stated price.[34] Options are often referred to as "derivatives" because their value derives from, and their price reflects, the value of some other ("underlying") asset (or, sometimes, liability). That asset may be a share of stock, but it could be anything—a car, a piece of real property, or a patent. The option *writer* is an offeror, who commits to selling (or buying) the asset at a given price for a given period. Instead of buying (or selling) that asset outright, however, an option holder instead owns the *right* to buy (or to sell) the asset at a given price.[35]

Consider someone who owns a *call* option,[36] giving her the right (but not the obligation) to buy one share of XYZ stock for $120 any time before

---

[34] For more detail on options, *see, e.g.*, STEWART BREALEY, STEWART MYERS & FRANKLIN ALLEN, PRINCIPLES OF CORPORATE FINANCE (13th ed., 2020). *See also* Bradford Cornell, *The Incentive to Sue: An Option-Pricing Approach*, 19 J. LEGAL STUD. 173 (1990) (using a variant of option theory to value litigation). Cornell's insight is that the option to abandon a lawsuit partway through (if discovery reveals that the claim is worth less than the plaintiff initially believed) constitutes an "embedded option" that increases the value of filing suit in the first place. *Id.* at 177. *See also* IAN AYRES, OPTIONAL LAW: THE STRUCTURE OF LEGAL ENTITLEMENTS (2005)(exploring the relevance of options to legal analysis and legal theory more generally).

[35] RESTATEMENT (SECOND) OF CONTRACTS § 25 (AM. L. INST. 1981)(defining an option contract as "[a] promise . . . [that] limits the promisor's power to revoke an offer." The owner of a "right to buy" is just the recipient of an (irrevocable) offer to sell at a given price.). At common law, what makes the offer *irrevocable* is that it is backed by consideration—the offeree has paid separately for the right to keep the offer open. *Cf.* Dickinson v. Dodds [1874] 2 Ch D 463 at 471–72 (explaining that a gratuitous promise to hold an offer open until a given time was not binding on the offeror because it had not been separately paid for and hence lacked consideration). *See also* U.C.C. § 2-205 (AM. L. INST & UNIF. L. COMM'N 1977) (allowing merchants to make binding commitments to keep offers open ("firm offers") *without* consideration under certain circumstances).

[36] If it is helpful, you can think of owning a call option as entitling you to "call the asset over to you" (buy it). If you do decide you want to buy it, the owner must sell, because they made you an (irrevocable) offer to do so.

July 30, 2024.[37] The $120 is known as the option's *exercise price*, which simply means that the underlying asset can be purchased (the option can be exercised) for that amount. The value of the option depends on the price of the underlying asset: if XYZ is currently trading at $200 per share, the call option is very valuable. The holder of the call can use her right to buy the option for $120 from the party who "wrote" or sold the option, and then turn around and sell it for its market price, pocketing the $80 difference.[38] If a share of XYZ is currently trading at $10, by contrast, the value of the call option is small—it will be worthless unless the share price rises above the exercise price. (Who would want to exercise their right to buy for $120 when they could easily buy for $10 on the open market?) The chances of that happening are presumably quite low.[39]

Conversely, consider the owner of a *put* option.[40] This gives the owner the right (but, again, not the obligation) to force someone else to buy the asset from them at the exercise price, which occurs when a counterparty has made an irrevocable offer to buy it at the exercise price. Someone who owns a put option on a share of XYZ stock with an exercise price of $120 (and an expiration date of July 30, 2024) could insist that their counterparty buy the share from them at that price. If XYZ is trading at $150 on July 30th, the right to sell a share for $120 is worth nothing, since the owner of the put option could always sell at the higher market price instead. But if a share is trading at, say, $100, then the owner of the put option can buy a share at the market price and then turn around and sell (that is, force their counterparty to buy) at the exercise price, pocketing the $20 difference. When the asset is worth less than the exercise price, a put option is a valuable thing to own.

---

[37] This is a so-called American option. A European option is exercisable only on the exercise date. But a famous observation in finance is that it's never (with provisos) worthwhile to exercise a call before its date; if you have reason to exercise it because the underlying asset's price is above the exercise price, you'll always do better selling the option instead.

[38] In practice, the option-holder would likely just settle-up for the $80 difference, reducing transaction costs.

[39] While the profit from holding an option is easily computed, the appropriate *price* to charge for an option is anything but. Fischer Black and Myron Scholes revolutionized finance (and Scholes won a Nobel prize) for deriving the correct formula for pricing an option. *See* Fischer Black and Myron Scholes, *The Pricing of Options and Corporate Liabilities*, 81 J. POL. ECON. 637, 640 (1973). BREALEY ET AL., *supra* note 34, at 573–79 (discussing the Black-Scholes formula in greater detail).

[40] If it is helpful, you can think of "shot *putting* the asset away from you" to someone else—forcing them to buy it at a given price. AYRES, *supra* note 34, at 205.

B. Position Diagrams

The careful reader will have discerned that the two kinds of options (call and put) and two kinds of positions one can take—write (commit to buy/sell) or own (purchase the right to buy/sell)—generate four basic financial stakes that can be created by these options.[41] Each kind of financial stake gives rise to a different relationship between the value of the underlying asset and the value of the option itself—the profit or loss that is realized from holding a given position. These so-called position diagrams are commonly used to depict the value of the underlying asset on the horizontal axis and the profit or loss of the option (holder or writer) on the vertical axis.

First consider the owner of a call option. To fix ideas, assume that the option is the right to buy a share of XYZ stock for $120, and that today is the option's expiration date. We want to graph the option holder's profit from owning the option as a function of the price of a share of XYZ stock. To simplify the graph a little, we will assume that the owner paid nothing for that option.[42]

If a share of XYZ has a market price of $0 the moment it expires, the option to buy it for $120 is itself worthless: why would anyone pay $120 for something they could buy for $0 on the open market? Of course, the same logic applies to *any* price below the $120 exercise price—the owner of the option would not choose to exercise it, so the profit from holding it would be . . . nothing at all. As the price of an XYZ share goes above the $120 exercise price, however, the call option-holder *would* want to exercise her option to buy. If the price were, say, $131, the holder could exercise her option, buy the share for $120, and then turn around and sell it for $131, turning an $11 profit. At an even higher price—say, $200—the option holder's profit would be $80 ($200 - $120) . Thus, the call option-holder's profit is zero for any price of the underlying asset less than the $120 exercise price; and that profit rises by $1 for every dollar that the price of the

---

[41] In practice, options are often combined in various ways, including owning the underlying asset on which the option is written, so there are many more possibilities than we explore here. *See supra* note 34 (outlining additional details on options).

[42] Of course, this is not only unrealistic, it likely runs into contract law problems discussed *supra* note 35 (promising to keep an offer open might not constitute a valid contract if it is not paid-for.). But the assumption makes the analysis a bit cleaner, so we'll stick with it. We will also assume throughout that the owner or writer of the option does not actually own the underlying asset involved. In a thick market, the commitment to sell can always be kept by purchasing the share on the market and then selling it.

underlying asset exceeds the exercise price.[43] Figure A1.A illustrates.

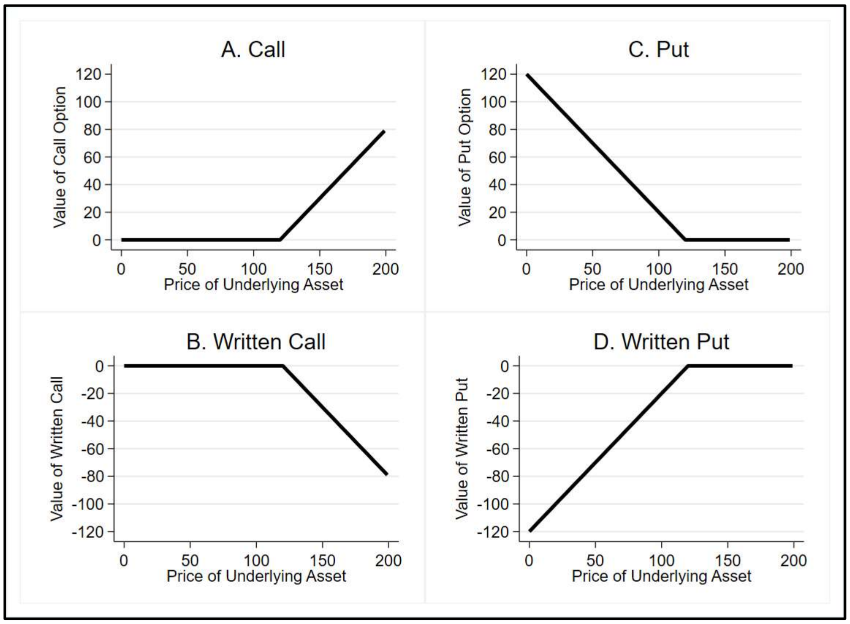

Figure A1: Value of Option Position as a Function of the Price of the Underlying Asset for Different Options (Exercise Price = $120)

What about the other side of this transaction? Suppose that instead of holding the option to buy a share of XYZ at $120, you had instead *written* that option, obliging you to sell at $120. In that case, a price of $0 for XYZ means that the option will surely not be exercised, and you will be out nothing. The same applies at any asset price below the exercise price. Once the price of a share exceeds $120, however, the option-holder will want to exercise it. You will then need to purchase a share of XYZ for its market price and then immediately sell that share to the option-holder for $120; you will of course be out the difference.[44] A graph of your profit looks like Figure A1.B. It is the mirror image of Figure A1.A, which makes sense because holding a call option has the opposite financial consequences of writing

[43] The formula for the option-holder's profit is Profit = *Max*[(Price - $120), 0], where "Max" means, "whichever is bigger." When the price of the underlying asset is less than the $120 exercise price, (Price - $120) is less than zero, the option won't be exercised, and the holder's profit is 0. When the price of the underlying asset is above $120, (Price - $120) is greater than zero, and that's the profit from exercising the option.

[44] Algebraically, your profit will be: *Min*[0, ($120 - Price of XYZ)], where "Min" means "whichever is smaller."

one—whatever the holder gains, the writer loses, so the two positions must always net to zero.

We now consider the other flavor of option, a put, which gives the owner the right to sell an asset at the exercise price.[45] If you own a put option on XYZ stock with an exercise price of $120, you can force someone else to buy it from you for that amount. If the market price of a share is $0, your option to make me pay you $120 for it is worth $120—you can make me pay you $120 for a worthless asset. As the market price of a share rises, the value of your put option declines, reaching $0 when the market price of the share hits the exercise price. When the market price of a share of XYZ exceeds the $120 exercise price, you will not want to exercise your option to make me buy it. That means the value of the put option is zero. The position diagram for the holder of a put option is shown in Figure A1.C.

Finally, consider the *writer* of a put option. The put writer's financial position is just the opposite of the holder of the put. When the price of the underlying asset is $0, the owner of the put option will want to force the writer to buy the asset for its exercise price ($120). That is a loss of $120 for the writer and a corresponding gain of that amount for the put holder. When the price of the underlying asset is greater than $120, the holder will decline to exercise the option, and the writer of the put loses nothing. This is illustrated in Figure A1.D.

C. HOMOLOGY WITH INSURANCE LITIGATION

It should now be clear that the insurer's position diagram in Figure A1.A is the equivalent of a put option with an exercise price of $120,000. (Since the put is on a liability—the award in the lawsuit—rather than an asset, the sign is negative, rather than positive.)[46]

Similarly, the insured policyholder in Figure A1.B has the equivalent of a written call option with an exercise price of $120. He pays nothing if the award (or settlement) is below the policy limit, but bears the expense of any payment above the limit.

---

[45] BREALEY ET AL., *supra* note 34.

[46] More generally, all insurance can be thought of as buying a put option. For example, if you have insured your house for $100,000, you have the right, but not the duty, to force the insurer to "buy" it from you for the $100,000 exercise price when certain conditions are met. If it burns down and the house is worth nothing, you gain the full $100,000 by forcing the insurer to "buy" it. If it is worth $25,000, you can still force the insurer to buy it for the exercise price. You would then get a check for $100,000 but give up a house worth $25,000, so you would net $75,000. In practice, of course, you wouldn't actually sell the house to the insurer: instead, you'd simply collect the $75,000 difference between the insured value and the actual value.