

PROPERTY

Shapley Analysis: A Cautionary Tale

Lloyd Shapley's pricing algorithm took the copyright world by storm, but does it really work?

◆ BY DOUGLAS LICHTMAN

In the United States, an administrative agency called the Copyright Royalty Board (CRB) sets rates for various compulsory licenses. Under those licenses, qualified parties can pay a government-set fee and then use implicated work without obtaining direct permission from relevant copyright holders. Some of the licenses allow cable television systems to retransmit, at regulated rates, copyrighted content originally aired on broadcast television. Others authorize companies like Pandora and Spotify to stream copyrighted music on their technology platforms, again without the need for direct negotiation.

The CRB has traditionally used a range of tools to set prices for these obligatory licenses. For instance, the CRB's three judges have historically used simulations to identify plausible rates. They have also used, as benchmarks, privately negotiated deals involving similar rights and similar parties.

But in 2015 the CRB took the first step in what would become a critical change to its decision-making processes. At issue was the final distribution of monies that had been deposited by cable system operators as legally required payment for the right to retransmit certain television programs that had already been broadcast on regular, over-the-air television. The CRB was responsible for distributing this money and, although the judges had already distributed approximately \$127 million to relevant parties, a residual \$1 million remained to be allocated between two final copyright claimants. The controlling statute did not dictate any particular standard for rendering this allocation. The judges, however, had previously committed to distribute funds according to the "relative marketplace value" of the respective programs, consistent with the "hypothetical market that would

DOUGLAS LICHTMAN is professor of law at the UCLA School of Law. This essay draws from written testimony he submitted to the Copyright Royalty Board in 2022 and from his article "Shapley Analysis—A Cautionary Tale," *Harvard Journal of Sports and Entertainment Law* 15(1).



exist but for the compulsory license regime.”

That framework could have led to very traditional types of economic analysis. The judges could have considered evidence from analogous markets, for instance, or relied on simulations, all while asking conventional questions about marginal cost and competitive entry. But they instead turned to a game theoretic construct known as the Shapley value, which they described in sweeping terms as “the optimal measure ... of relative value in a distribution proceeding.”

SHAPLEY'S ALGORITHM

What is this magical approach? The Shapley construct is in essence an algorithm for dividing economic returns in instances where some number of distinct entities together generate a shared profit or incur a shared cost. It is said to achieve a “fair” allocation of that benefit or burden between the relevant parties by accounting for each party’s marginal contribution to the whole. Mathematician and economist Lloyd Shapley, who spent most of his career at UCLA, first sketched the idea in a 1951

paper that he wrote while still a grad student. A revised version was published in 1953; and, in 2012, Shapley won the Nobel Memorial Prize in Economics in part for this work.

The idea is typically expressed in mathematically sophisticated ways, but the concept can be introduced using a simple example. Imagine that three friends are leaving a restaurant to travel to their respective homes. They can each take different taxis, but because their paths overlap, they decide to hire one taxi and share the cost. The most efficient, straight-line route is to drop Ann first, then Bob, then Chris, and the friends expect the meter will show \$30 when arriving at Ann’s home, \$44 when arriving at Bob’s, and \$54 when arriving at Chris’s. The question for the friends is how to divide the total \$54 fare among them, given the various partial overlaps.

One option would be to focus on the actual order in which the friends arrive home. Under this approach, Ann would pay the initial \$30, because that is the fare associated with her part of the trip. Bob would pay the next \$14, which is the additional cost required to travel from Ann’s house to Bob’s. Chris would then pay the final \$10, as he at that point is riding alone. This allocation has the virtue



PROPERTY

of being administratively simple: the riders each pay the amounts due whenever they exit the cab. But the allocation disproportionately favors Chris. After all, Chris literally enjoys a free ride for the entire shared portion of the trip; he contributes only to the final portion, a portion that exclusively benefits him anyway.

Consider, then, an alternative “arrival” sequence, such as perhaps Bob, then Chris, then Ann. Under this pattern, Bob would pay \$44, which is the total cost for his part of the ride; Chris would pay \$10 when he exits; and Ann this time would be the lucky one, because in this pattern the full fare is covered without her paying a dime. Note that the “arrival” concept here is conceptual: the taxi would still take the same path it did in the prior one, but the payment obligations change based on the newly proposed theoretical order.

Shapley’s algorithm balances these various scenarios by cataloging every possible permutation, calculating the resulting payment patterns, and ultimately averaging the payments to determine each person’s “fair” share. In this example, there are six possible sequences to consider, which we can identify using the riders’ first initials: ABC, ACB, BAC, BCA, CAB, and CBA. Those options, the fares, and the resulting averages are shown in Table 1. The upshot is that, out of the total \$54 owed, Shapley’s approach would have Ann pay \$10, Bob \$17, and Chris \$27. And, as promised, that result does align with an intuitively “fair” outcome. The first leg of the trip benefits all three riders, so they each pay one third of the \$30 cost. The second leg benefits just Bob and Chris, so they each pick up half of that next \$14. The final leg benefits Chris exclusively and so he pays the final \$10 himself. Ann ends up paying a total of \$10, Bob a total of $\$10 + \$7 = \$17$, and Chris $\$10 + \$7 + \$10 = \27 .

Again, Shapley’s paper introduced all this more formally, and with none of the public policy overlay. He defined his model mathematically and, instead of offering a concrete example like individual riders sharing portions of a journey, he articulated the formation of abstract “coalitions” where “players” team up to generate unspecified economic returns. As to the policy overtones, Shapley in his paper did not articulate any specific notion of fair play, nor did he champion any specific applications for the algorithm. Instead, perhaps because he was writing at a time when game theory was

still a relatively new field of inquiry, his focus was on the nuts and bolts of the modeling. Indeed, the bulk of his paper was invested in proving certain axioms about his approach, including that the sum of the payments add up to no more, and no less, than the actual total (confusingly, he called this there-is-no-waste property “efficiency”), and that his approach yields consistent results if, instead of considering the full interaction in the context of a single model, a modeler were to break the interaction into smaller subgames, analyze those, and then combine the payoffs.

THE CRB ENDORSEMENT

The CRB said little in the 2015 order that first praised Shapley analysis as the “optimal” and “most efficient” approach. The parties to that particular proceeding had not themselves proposed Shapley analysis, and thus they had not submitted any of the necessary evidence or testimony. But the CRB’s enthusiasm was clear, and stakeholders reacted to it.

The next year, University of Toronto economist Joshua Gans filed a report with the CRB applying Shapley analysis in the context of a proceeding related to online music streaming. Then came a report from Duke University economist Leslie Marx, then one from the University of Canterbury’s Richard Watt, and then one from Princeton’s Robert Willig, all applying the Shapley construct to various licenses under the CRB’s regulatory purview. The CRB, meanwhile, continued to endorse the approach. Rate decisions published in 2019 and 2023 explicitly relied on Shapley analysis to derive “reasonable rates and terms” for a specific license that governs the use of a song’s words and notes. And, in an intervening appellate case, the judges defended their approach before the DC Circuit, convincing a unanimous appellate panel that this approach to pricing fell “well within the Board’s discretion.” No party, meanwhile, meaningfully challenged the relevance or reliability of Shapley analysis, as nearly every stakeholder had sponsored an economic expert who had praised it or at least tolerated it.

REJECTING SHAPLEY ANALYSIS

Copyright licensing could have been left entirely to the unregulated market. Copyright holders in this scenario would have negotiated directly with would-be copyright licensees. Traditional market forces would have defined the necessary terms and shaped the necessary rates. Congress created a system of compulsory licenses, however; and from that, a decade of testifying economists have all reasonably inferred that the CRB is supposed to do something more than simply recreate market outcomes. Yes, these economists have looked to the market for information about plausible rates, incentives, and behavior. But they each have then urged the CRB to deviate from those actual or hypothetical market results to address one or another specific market imperfection.

Some of the economists worried about market power. Marx, for instance, championed adjustments to offset what she perceived to be the undue leverage enjoyed by copyright holders because of concentration in the music industry. Others worried about

Table 1
A Shapley Analysis Example

How much should each taxi passenger pay?

Conceptual Payment Order	Amount Paid		
	A	B	C
A, then B, then C	\$30	\$14	\$10
A, then C, then B	\$30	\$0	\$24
B, then A, then C	\$0	\$44	\$10
B, then C, then A	\$0	\$44	\$10
C, then A, then B	\$0	\$0	\$54
C, then B, then A	\$0	\$0	\$54
Averages:	\$10	\$17	\$27

imperfect information. Watt, in this spirit, argued that copyright holders face a considerable challenge when negotiating with firms like Apple or Amazon because, to an unknown degree, those firms use streaming to drive business to other products and services.

These are potentially valid concerns, in my view, and they warrant thoughtful exploration. But Shapley analysis does not speak to any of them. Start with the claim that it can be used to model “market” interactions. This is a foundational, explicit, descriptive claim for Gans, Watt, and Willig, and they each use it to justify Shapley analysis as a framework for answering all the other open questions. Gans, for instance, asserted in his original 2016 filing that Shapley analysis can be used to estimate the royalties “that would prevail in an unconstrained *market*” (my emphasis, here and below). Watt wrote in 2017 that Shapley analysis “mimics what a

any plausible market, after all, one of the two right-glove proprietors would consummate the deal and earn a return, while the other would strike no deal and earn nothing. That risk would, in turn, play a critical role in the real interaction, in that it would motivate competition between the two right-glove sellers, with each trying to undercut the other’s price for fear of otherwise ending the interaction empty-handed (sorry). But there is no real-world market where, after one transaction, both right-glove sellers are nevertheless paid. And there is certainly no real-world market where the successful right-glove seller and the unsuccessful right-glove competitor earn the exact same return.

Gans nevertheless leapt from that implausible example to the real issue that was then before the CRB. He acknowledged that the “usual intuition” is that competing parties “can be played

off against one another to effectively be pushed to receiving payments close to their costs, earning no surplus”; but he asserted that the Shapley value approach “predicts” (predicts?) instead an outcome where all competitors are paid. He then offered a just-so story that could have been used to justify almost any numbers he might have proposed. His story focused on the potential for “left-glove” copyright holders to pit “right-glove” technology companies against one another in a bidding war. He

announced that copyright holders would do no such thing. Because streaming services “have a role in providing competition against one another,” he explained, copyright holders “will not push these streamers to their limits in negotiation.” Instead, according to Gans, copyright holders will leave precisely the Shapley value on the table, using that exact amount to strategically keep one streamer “waiting in the wings” as a competitive check on the other.

Marx filed the next report to seriously consider Shapley analysis, and she similarly implied that it can do much more than identify a “fair” allocation of some specific shared gain or loss. Marx was concerned that copyright holders might have “concentrated market power” given industry consolidation. She knew that a simple Shapley model would do nothing to mitigate that distortion, so she presented a model where she “intentionally elevated the market power of the [streaming] services” by using a single player to represent all interactive streamers rather than modeling each existing streamer separately. As she explained in live testimony, her intent was to offset copyright holders’ market power by introducing market power on the other side of the transaction. But that was another just-so adjustment, one that shifted the numbers in the desired direction, sure, but did so to a completely arbitrary degree. Shapley mathematics offered no insight into the extent of the original market distortion. The Shapley dynamic neither suggested nor validated Marx’s attempt to offset it. And of course not, because Shapley models are not models of market behavior.

Shapley mathematics offers no insight into the extent of the original market distortion. And of course not, because Shapley models are not models of market behavior.

free and unrestricted *market negotiation* would yield” and followed up in 2021 with the assertion that the methodology “reflects effective *competition*.” Willig similarly assured the judges in 2019 that “Shapley Values are an appropriate approach for assessing rates that would be negotiated in the hypothetical *marketplace*.”

But how can these descriptions possibly be true? Consider Gans’s 2016 filing. He was the first economist to take seriously the CRB’s suggestion that Shapley analysis be used as a framework for rate-setting, so he understandably opened the relevant portion of his report with an example designed to teach the basic operation of Shapley mathematics. His example involved three firms selling gloves. Two each offered a single right-hand glove, while the third offered a single left glove. The surplus generated from matching the left glove with one of the right gloves was defined to be \$1, and there was no value associated with an unpaired glove. Gans articulated the familiar Shapley process where all possible coalitions of glove providers “arrive” in all possible orders. He concluded that each provider of a right glove should be assigned a Shapley value of $\$1/6$, whereas the lone provider of the left glove should be accorded a higher Shapley value of $\$2/3$. Gans explained that the provider of the lone left glove “commands a higher share of the surplus because she is the only player to own a left glove,” whereas the two providers of potential right gloves “are substitutes for one another” and hence compete away some of the value that a lone right glove owner would otherwise receive.

This model bears no resemblance to any real-world market. In

PROPERTY

**MODEL AMBIGUITY AND
STAKEHOLDER INCENTIVES**

The taxi example makes these problems even more plain. Admittedly, that example is a significant simplification of Shapley's original, sophisticated model, and it pales, too, in comparison to the mathematical extensions that have been developed since. But simplicity lays bare the actual workings of a model, and here simplicity reveals two fundamental elements that make Shapley analysis plainly inappropriate for the CRB's purposes.

First, Shapley's structure leaves no room for ambiguity as to how many and which specific parties ought to be considered legitimate stakeholders. Three riders are relevant to the taxi example. There is no mechanism by which to explore whether Ann and Bob should be counted as a single passenger because they are dating, or whether Chris should count double because he is bringing along heavy luggage. There are three riders, that fact leads to a chart with six possible payment orders, and the addition or subtraction of even one rider would significantly alter every calculation.

Second, Shapley analysis is unapologetically static, with no room for players to engage in strategic behavior and no accounting for the long-run incentives created by the model's proposed allocations. Chris in the taxi example cannot negotiate a better deal by credibly threatening to ride alone. Alice and Bob cannot tweak their allocations even if they realize that, at these numbers, Chris will next time choose a restaurant closer to his home or opt to drive his own midlife-crisis Ford Mustang.

Note that these are not criticisms of the Shapley approach per se. Quite the opposite; Shapley analysis largely resonates in the taxi example in that friends often find themselves in interactions where the number and identity of the participants is given and strategic play is unlikely because it would violate powerful social norms. That is, when friends share a taxi, split a restaurant bill, or—Marx's intuitive example—share a boat ride, they very plausibly are looking for a static “fair” outcome by which they will then non-strategically abide. In government rate-setting, by contrast, none of that holds true.

Competitors, not friends / Consider in this light the definitional questions about who the relevant stakeholders are, how many of them will share in any allocation, and thus implicitly what monies ought to be deemed eligible for division. Again, those questions all have obvious answers in the taxi example. Amy, Bob, and Chris are the only riders. The total taxi fare is the only number in play. At the CRB, by contrast, these same questions were the subject of real and plausible dispute. Gans, Watt, Marx, and Willig vigorously disagreed about whether copyright holders should be represented in their various models as a single unified rightsholder, as one representative record label and one representative music publisher, or as some larger number of separate

players each representing (in those proceedings) a real-world record label, a real-world music publisher, and possibly even a real-world singer, musician, producer, or songwriter. Gans, Watt, Marx, and Willig disagreed, too, on the question of how best to represent potential copyright licensees. Marx, for instance, argued that an appropriate Shapley model would include not just some number of players representing the streaming services, but also some number of additional players standing in for other types of distribution partners who also contribute to the overall market for music. Her intuitive point was that a “fair” allocation of copyright royalties can only be made by considering all the ways the implicated copyrights and the implicated streaming technologies interrelate. Watt thought this approach flawed,

Shapley analysis is unapologetically static, with no room for players to engage in strategic behavior and no accounting for long-run incentives created by the model.

agreeing that other types of music monetization are relevant but asserting that substitution and promotion across platforms should be measured in other ways.

Whatever the right answer, these are critical inputs to Shapley analysis in that they significantly affect the math. Consider a Shapley model where a painter, a decorator, and a furniture maker can potentially team together to modernize an apartment. For simplicity, ignore costs. If modernization generates \$120 in value but can only be accomplished through the combined efforts of all three players, the Shapley procedure will allocate \$40 in value to each. Redefine the model to require a fourth necessary player—such as a real estate agent to market the finished apartment—and in response the Shapley algorithm will reduce the payments to \$30 per player. Redefine the model again such that the painter is newly conceptualized as a lead painter and two assistant painters, all necessary, and now the Shapley values drop to \$20 for each painter, \$20 for the decorator, \$20 for the real estate agent, and \$20 for the furniture maker. Make another change—for example, framing the model such that any one of the three painters can do the entire job alone—and the Shapley values again shift considerably, this time with each painter being accorded \$10 while every other skilled contributor earns \$30. Shapley models, in short, are extremely sensitive to the assumed number and types of players included. And at the CRB, in sharp contrast to the taxi example, those values are significantly vulnerable to both strategic advocacy and genuine dispute.

Just to be clear, my concern here is neither the generic concern that a model's inputs drive its outputs, nor the generic concern

that models inevitably must approximate reality rather than completely capture it. My concern is that Shapley models are particularly sensitive to their inputs, and hence this modeling approach is unreliable when those inputs are (at best) significantly stylized simplifications. So, while a Shapley model might have much to teach when the parties being modeled are the members of a single family or voting members of a governmental institution like a court or legislature—that was the first practical application Shapley himself pursued—the Shapley approach is significantly less reliable where, as here, the real-world cast is much larger, much more diverse, and in countless ways intertwined.

Turn next to the even more problematic point, that Shapley's model is completely static. Copyright law is an incentive system, recognizing in authors certain exclusive rights in order “to promote the progress of Science and useful Arts.” That process is intentionally dynamic. The whole idea is to inspire strategic responses from (say) singers, songwriters, musicians, producers, record labels, music publishers, and, yes, streaming services, technology companies, and listeners, too. All these stakeholders are supposed to calibrate their actions in response to the returns they expect to receive, the fees they expect to incur, and the rights and privileges they otherwise expect to enjoy. To use an allocation mechanism that fully ignores dynamic implications is to measure a three-dimensional space using a two-dimensional tool. Allocations cannot plausibly meet the statutory requirements of being “fair” and “reasonable”—let alone efficient or consistent with any plausible legislative purpose—if they are being made while blind to the bigger dynamics that are core to the underlying legal rule.

Further, even if it were somehow appropriate to allocate copyright monies without regard to long-run incentives, static analysis would still be inappropriate in this context because, while CRB analysis might be static, the copyright marketplace is not. In its 2019 and 2023 decisions, for example, the judges used Shapley analysis to establish what was intended to be a “fair” allocation of monies between the interactive streaming services, the record labels, and the music publishers. The CRB's ruling was based on its view as to how much money each stakeholder ought to retain. But the proceeding itself established only the rate due music publishers for one specific copyright right. The ruling did not constrain what record labels could charge for copyright rights they control, for instance, and indeed the CRB had no power to do so even had it tried. Record labels were thus immediately free to react to the CRB's rate by charging whatever license fees the market would bear, even if that number was higher than the CRB's calculated amount, and even if paying it meant that streamers would end up with less money than the CRB intended. For this reason, too, Shapley analysis falters in this application. If the Shapley value assigned to music publishers is appropriate at all, it is appropriate conditional on record labels and streaming services also being accorded their Shapley values. In the real world, however—because in any given proceeding the CRB has power over only a subset of the relevant rights and a subset of the relevant parties—that condition dramatically fails.

THE STAKES

Rate-making proceedings at the CRB are anything but welcoming. The typical proceeding lasts years and generates literally thousands of pages of evidence. Testifying expert economists, meanwhile, seem to speak in tongues. Even the judges, from time to time, complain that years into any given proceeding they are unsure as to what certain stakeholders are arguing or what assumptions are being made in support of competing rate proposals. Combine that with a 70-plus-year-old mathematical construct and there is a very powerful case to be made for ambivalence. The show is not worth the price of admission. The arcane details of Shapley analysis and the minutiae of CRB ratemaking are best left to insiders alone.

But the rates at issue in these proceedings matter. In the music industry, for instance, compact disc sales have plummeted over the past 20 years, with consumers spending \$13.2 billion on the format in 2000 but only \$483 million in 2022. Direct sales of digital singles and albums do not come close to filling the gap, amounting to barely \$465 million in 2022. And, while the biggest stars might be able to earn substantial sums by touring (hi, Taylor) or by licensing their music for use on television and in movies, those options are unavailable to the overwhelming majority of singers and songwriters, who cannot fill stadiums and whose songs will never be picked for those types of use.

Billions of dollars, by contrast, are at stake every year in transactions governed by the CRB, with Amazon, Apple, Spotify, Google, Pandora, and iHeart each either paying CRB-set rates or negotiating private licenses in their shadow. This is the money that will drive the music industry in the foreseeable future, and the money that will similarly drive the development of streaming and other technological advancements. And, while there is no easy formula for allocating those funds, my point is simply that Shapley analysis—the current belle of the ball—in fact offers no helpful insight. Shapley models are not descriptive of real-world markets and hence they do not help the CRB characterize the free-market baseline. Shapley models offer no tools by which to measure market imperfections, even though the central reason to have compulsory licenses is to account for those failures. And, worst of all, Shapley modeling is unapologetically static: a puzzling mismatch for proceedings relevant to copyright law, a set of rules fundamentally designed to inspire strategic responses and create long-run incentives.

In short, this application of Shapley analysis is a mistake, and it is time for the CRB to abandon the approach. R

READINGS

- “A Value for n-Person Games,” by Lloyd S. Shapley. In *Contributions to the Theory of Games*, Vol. 2, edited by Harold W. Kuhn and Albert W. Tucker; Princeton University Press, 1953.
- “Fair Copyright Remuneration: The Case of Music Radio,” by Richard Watt. *Review of Economic Research on Copyright Issues* 7(2): 21–37.
- “Getting Pricing Right on Digital Music Copyright,” by Joshua S. Gans. *Review of Economic Research on Copyright Issues* 15(2): 1–22.