

DISTRIBUTING ATTORNEY FEES IN MULTIDISTRICT LITIGATION

Edward K. Cheng *, Paul H. Edelman[†] and Brian T. Fitzpatrick[‡]

ABSTRACT

As consolidated multidistrict litigation has come to dominate the federal civil docket, the problem of how to divide attorney fees among participating firms has become the source of frequent and protracted litigation. For example, in the National Football League (NFL) Concussion Litigation, the judge awarded the plaintiff attorneys over \$100 million in fees, but the division of those fees among the twenty-six firms involved sparked two additional years of litigation. We explore solutions to this fee division problem, drawing insights from the economics, game theory, and industrial organization literatures. Ultimately, we propose a novel division method based on peer reports. Participating firms assess the relative contribution of other firms to the litigation, and then optimization or Bayesian techniques arrive at a consensus or compromise fee allocation. Our methods are intuitively easy to understand, enable broad participation, and are resistant to collusion or other strategic behavior, making them likely to be accepted by the firms involved. We thus provide courts with an important mediation tool or decision rule for these fee division disputes.

* Hess Professor of Law, Vanderbilt University, Nashville, TN, USA; Visiting Professor (2020), Hebrew University of Jerusalem, Jerusalem, Israel; Vanderbilt Law School, 131 21st Avenue South, Nashville, TN 37203, USA. Email: edward.cheng@vanderbilt.edu

[†] Professor of Mathematics and Professor of Law, Vanderbilt University, Nashville, TN, USA.

[‡] Milton R. Underwood Chair in Free Enterprise and Professor of Law, Vanderbilt University, Nashville, TN, USA.

We are grateful to Ben Alarie, Roy Baharad, Lynn Baker, John Cunningham, Bert Huang, David Madigan, Rahul Mazumder, Linda Mullenix, Liam Paninski, Charlie Silver, Albert Yoon, two anonymous referees, and the editors for helpful comments on this article, as well as to the participants at workshops and classes at George Mason University, Harvard University, the University of Texas, the University of Toronto, and Vanderbilt University. Thanks also to Brooke Bowerman and Ryli Wallace for superb research assistance. Brian Fitzpatrick has served as an expert witness in attorney fees allocation disputes, including the NFL Concussion Litigation.

1. INTRODUCTION

In 2015, the NFL and its retired players reached a widely reported final settlement in the NFL Concussion Litigation (Belson 2015).¹ The (uncapped) settlement agreed to payments up to \$5 million for former players, and as of January 2021 involved over 3,000 submitted claims and over \$800 million in payouts (NFL Settlement Program Summary Report 2021). It was, in many ways, representative of modern complex litigation. It featured large numbers of plaintiffs and attorneys, massive dollar amounts, and of course, settlement.

Widely unnoticed and hidden behind the headlines, however, was another complicated and protracted dispute—this one over attorney fees. After the settlement was affirmed on appeal, the trial judge decided the lawyers should be paid \$112.5 million, but the case involved a multidistrict consolidation involving twenty-six firms. The judge thus had to allocate the fees among the twenty-six law firms. She asked the “lead counsel” of the multidistrict litigation (MDL) to propose an allocation.² Lead counsel in turn divided the money based on an “adjusted lodestar” method, weighting the number of hours worked by a multiplier to reflect each firm’s importance to the litigation. Lead counsel’s own multiplier was almost four times the base rate, and he proposed that 60 percent of the \$112.5 million fee award be awarded to his own firm.³

Unsurprisingly, almost half of the other firms filed formal objections to the proposed allocation. Some proposed using a neutral third party like a special master. The judge instead made some minor changes and awarded the fees as proposed. Dissatisfied attorneys then appealed the allocation to the Third Circuit, which eventually affirmed the allocation after another two years of litigation. All told, the dispute over attorney fees took almost as much time as the lawyers took to secure the original ~\$1 billion settlement from the NFL.

The attorney fee problem we see in the NFL Concussion Litigation is far from unique. At the end of most consolidated MDLs, the court must decide how to pay the group of lawyers that prosecuted the case on behalf of plaintiffs. MDL attorneys are not selected and monitored by clients like other lawyers, so there are no *ex ante* contractual arrangements governing fees. Courts therefore have to award fees *ex post*, and they must do so in ways that provide incentives

1 See also *In re Nat'l Football League Players' Concussion Injury Litig.*, 307 F.R.D. 351, 361 (E.D. Pa. 2015), amended sub nom. *In re Nat'l Football League Players' Concussion Injury Litig.*, No. 2:12-MD-02323-AB, 2015 WL 12827803 (E.D. Pa. May 8, 2015), aff'd sub nom. *In re Nat'l Football League Players Concussion Injury Litig.*, 821 F.3d 410 (3d Cir. 2016), amended (May 2, 2016).

2 The facts of the attorney fee dispute are recounted in *In re NFL Players' Concussion Injury Litig.*, No. 2:12-md-02323-AB; MDL No. 2323, 2018 U.S. Dist. LEXIS 88055 (E.D. Pa. May 24, 2018).

3 One of us (Fitzpatrick) filed an expert declaration in support of the allocation.

for good work and that are sufficiently fair to avoid acrimony and further litigation.

This is not a small matter. Over half of all civil litigation pending in federal courts right now has been rolled into MDLs (Resnik 2017, p. 1771 and figure 3). Lacking an established procedure for awarding fees, judges are forced to use ad hoc compensation schemes, and these schemes are often deeply controversial. The result is the possibility of suboptimal incentives and prolonged, expensive, and inefficient satellite litigation over fees. Bringing order to this aspect of our civil justice system is long overdue.

Conceptually, the process of determining attorney fees consists of two fundamental issues. The first is how much to pay the attorneys in aggregate, and on this question, there has fortunately been considerable legal scholarship (Clermont & Currivan 1977; Rubenstein 2009; Eisenberg & Miller 2010; Fitzpatrick 2010a,b). The second is how to divide the pool amongst the group, but unfortunately on this second question there is little guidance from the literature. In practice, the lead lawyer often divides the fees, although sometimes the court chooses to do so itself or steps in when disputes arise.⁴ But *how* should the fees be allocated? Relying on a court's subjective impressions to do the division is clearly suboptimal. For one thing, mass litigation is complicated, and courts will frequently lack detailed, first-hand knowledge of each participant's contribution to the litigation. For another, these impressions are vulnerable to charges of favoritism, bias, and error. They also encourage subsequent challenges and are difficult to review on appeal.

Just as in the NFL Concussion Litigation, fee allocation thus often falls back on a familiar method: the lodestar. The problem with the lodestar method is that no one likes it. Its principal vice is well known: if lawyers are paid for the hours they work, they have no incentive to work efficiently, or worse yet, they have incentives to exaggerate the hours reported (Fitzpatrick 2010b, pp. 2051–2052). These vices are especially acute in the MDL context, which lacks sophisticated clients to monitor and supervise the lawyers. The multiplier in an adjusted lodestar may blunt this problem, but only at the risk of resentment. The other firms in the NFL litigation were unhappy when the lead lawyer valued his own firm's work at nearly four times the base.

4 For examples in which courts did the allocation themselves, see *Allapattah Servs., Inc. v. Exxon Corp.*, 454 F. Supp. 2d 1185 (S.D. Fla. 2006) and *In re TFT-LCD (Flat Panel) Antitrust Litig.*, No. M 07-1827 SI.MDL. No. 1827, 2013 WL 1365900, at *7–16 (N.D. Cal. April 3, 2013). For instances, in which the lead lawyer did the allocation with review on objection by the court for abuse of discretion, see *Victor v. Argent Classic Convertible Arbitrage Fund L.P.*, 623 F.3d 82 (2d Cir. 2010), *In re Initial Public Offering Sec. Litig.*, No. 21 MC 92(SAS), 2011 WL 2732563 (S.D.N.Y. July 8, 2011), and *In re Vitamins Antitrust Litig.*, 398 F. Supp. 2d 209 (D.C. Cir. 2005).

In this article, we examine this problem of fee distribution. Part 2 sets the stage by reviewing how much to pay attorneys in mass litigation as a group. As we discuss, prior to a wave of critical scholarship in the 1980s, courts routinely used the lodestar method to award aggregate fees in successful mass settlements. Today, however, courts have abandoned the lodestar method, favoring the so-called “percentage method” instead. This percentage method creates the correct starting pool for MDL attorney fees, but does not tell us how to divide it among the group of attorneys.

Part 3 turns to the problem of distributing the fee pools created by the percentage method. We review the problems of using the lodestar method for division, and then explore some innovative ways to divide fees, including ones relying on artificial market mechanisms.

Finally, in Parts 4 and 5, we offer a novel solution based on peer assessment. Each firm rates the relative contribution made by its peers,⁵ and then our method uses this information to reconstruct what is effectively a “consensus” division. Our proposed method has several attributes that make it highly desirable as a solution to the division problem. First, relying on peer consensus promotes acquiescence and reduces the chance of disputes. Second, our method permits partial peer assessments. Not every firm in an MDL will be familiar with the work of all the other firms, so allowing firms to rate only those with which they worked improves information quality. Third, it resists collusive behavior, dampening the ability of colluding firms to inflate each other’s shares. And finally, our proposed method is conceptually easy to understand, hopefully making it attractive to courts and attorneys alike.

2. BACKGROUND

MDL consolidates what are often thousands of cases that share common issues.⁶ Each of these cases will already have its own lawyer, meaning that an MDL can involve hundreds if not thousands of lawyers. The only practical way to manage representation is for the judge to assemble a subset of lawyers—in other words, a team—to litigate the MDL issues.⁷ The team is often one or two dozen lawyers strong and organized by the judge into various committees, each of which will be responsible for one part of the litigation. For example, several committees might be responsible for drafting and opposing the various

5 In our initial proposal, no firm reports on its own allocation, although we later relax this restriction.

6 For data on the distribution of cases in MDLs, see <https://www.uscourts.gov/statistics-reports/judicial-panel-multidistrict-litigation-judicial-business-2018>.

7 A team can also become class counsel if a class action grows out of the MDL.

dispositive and nondispositive motions that will arise; another committee might oversee discovery; another might negotiate settlement; and another might prepare for so-called bellwether trials. The committees themselves will do much of this work, but some of it might be outsourced to new lawyers not part of the original team. All of these committees are usually overseen by yet another committee, often called an “executive” or “steering” committee. The lawyers on this oversight committee will have the most knowledge about what everyone else has done, but even they will be far removed from the details about how that work was done; the other lawyers on the various committees on the team will often be familiar only with the work that was done in their own separate fiefdoms (Klonoff 2000, pp. 173–196).

As a consequence, the MDL context lacks the usual mechanisms that incentivize attorney behavior and determine their fees. The conventional structure, in which clients hire, pay, and monitor attorneys breaks down because clients are effectively absent.⁸ Courts select the lawyers, and courts are poorly positioned to negotiate terms or monitor attorney performance. Courts have neither the time nor the interest, and are traditionally passive institutions (Fitzpatrick 2009).

In theory, the lawyers themselves could work out an *ex ante* arrangement at the beginning of the case. Indeed, given the high stakes, there are considerable incentives for them to do so. Yet, in practice, *ex ante* contracting is rare. One reason may be that since the legal team is assembled by the court, negotiating allocations may seem presumptuous. Another reason may be the uncertainty over negotiating partners, as the lawyer team often evolves as outsiders are brought in at different points in the litigation. Still, another reason is that judges often do not feel bound by such agreements.⁹ Perhaps with changes in norms or appropriate institutional reforms, we will see more *ex ante* fee division arrangements, but for now, *ex post* division is the reality we face.¹⁰

MDL attorneys thus present a classic agency cost problem. How can we ensure that they do what their (absent) clients would want them to do? For many decades now, scholars have recommended reliance on the invisible hand of incentives (Fitzpatrick 2010b).

The law and economics literature on attorney compensation is vast,¹¹ but for our purposes here, a few highlights suffice. To begin, the literature suggests

8 This breakdown also occurs in other contexts in which courts appoint attorneys, such as for minors or incapacitated persons.

9 See, for example, *In re Syngenta AG MIR 162 Corn Litig.*, No. 14-md-2591 (D. Kan., December 31, 2018) at 14 (“[T]he Court is not bound by a private agreement among attorneys.”).

10 We leave an exploration of mechanisms for encouraging *ex ante* bargaining to future work.

11 Much of this literature is summarized in Fitzpatrick (2019).

that in the missing-client context, paying attorneys a percentage of the recovery is superior to paying attorneys by the hour. Paying by an hour creates a perverse incentive to work inefficiently and provides no incentive to maximize client recovery. In contrast, the percentage method creates strong incentives to both work efficiently and maximize recovery. To be sure, the solution is not perfect. Since the lawyers bear all of the incremental labor costs but receive only some of its fruits, lawyers have incentives to settle prematurely—when *their* net recovery is maximized, not the client's (Epstein 2002; Shavell 2004). Various solutions to the premature settlement problem have been proposed, though none completely solve the problem in a practical way. One solution is to increase the lawyer's percentage as the recovery is improved (Coffee 1986). The leading theory involves a combination of percentage recovery and hourly rate (Clermont & Currivan 1977), but even that does not completely solve the problem without a complicated third-party insurance scheme (Polinsky & Rubinfeld 2003). On the whole, however, the percentage method remains preferable.

Much of what judges now do with respect to *total* attorney fees is consistent with this insight. Today, courts typically pay mass litigation lawyers (as a group) a percentage of what they recover—usually 6–12 percent in MDL, where the lawyers also eventually collect a percentage from their individual clients (Rubenstein & Newberg 2011; Burch 2019).

All of this is well and good. But it all assumes the presence of only one lawyer (or firm). That limitation is fine in the class action context, where judges typically appoint a single firm as class counsel (and where most of the literature has been directed). But MDL is different, because as noted, MDLs involve *teams* of lawyers. How should we award attorney fees when an absent client is represented by a team of otherwise independent lawyers? The bad news is that the percentage method is difficult to implement, and this is what gave rise to the problems we observed in the NFL Concussion Litigation. With the lodestar method, no such problem exists, because lodestar solves the problem automatically—attorneys receive fees based on hours worked. But with the percentage method, there is only a single common pool. Which lawyer should receive what portion of the common pool? If every lawyer received the same percentage, what incentive would any of them have to work? And if lawyers receive percentages based on their contribution to the outcome, how does one figure out those contributions?

Perhaps the lawyers can amicably work out a division among themselves, but that is unlikely since the lawyers will inevitably have biased perspectives about their individual worth. At the same time, courts will rarely have sound, independent information regarding which lawyer did what. Thus, in an MDL

involving multiple law firms, almost everyone falls back on the lodestar method for dividing the pool despite its obvious vices (Becker, Specter, & Kline 2016).

This state of affairs is a massive oversight. More than half of all federal civil litigation is wrapped up in MDL, and MDL will almost always involve multiple-lawyer teams.¹² These teams will lack pre-existing relationships or contractual agreements, yet they ultimately must divide the spoils.

The existing literature thus only addresses half the problem. It optimizes how to pay MDL attorneys in aggregate, but not how to pay the attorneys individually. It helps determine the pool, but not the division. In the Parts that follow, we tackle this division problem, extending the insights of the law-and-economics literature on the single-principal-single-agent problem into the single-principal-*multiple*-agent context.

3. THE FEE DIVISION PROBLEM

In this part, we first situate the fee division problem in the theoretical literature. After setting this theoretical framework, we examine some previously proposed methods that might address the problem.

3.1 Theory

The fee division problem involves several related conceptual goals. The first goal is equity. A division rule that lacks fairness, or at least *perceived* fairness, will fail to gain attorney acceptance and will generate wasteful secondary litigation. This problem of fair division—how a group of claimants can equitably allocate a good among themselves—has a long and storied history. Take, for example, the classic cake-cutting problem, which dates back almost three millennia to Hesiod’s *Theogeny* (Brams & Taylor 1996). To divide a piece of cake equally and foster acceptance of the result: “I cut, and then you choose.” The literature on fair division, both academic and popular, is extensive (Moulin 1988; Brams & Taylor 1996; Robertson & Webb 1998). For example, recent work in the area has extended the cake-cutting problem to instances in which the thing divided is not uniform, so that some parts have negative value or that different actors value various parts differently (Segal-Halevi 2018). These strategies are especially useful in dividing inheritance estates.

The fair division literature, however, primarily focuses on cases in which each claimant’s right is given. Either all of the parties are entitled to equal

12 Under the federal Multidistrict Litigation Act, any time more than one case presenting the same factual question is filed in a federal court, a special panel of federal judges can decide to consolidate and transfer all such cases to a single federal judge for pretrial purposes. 28 U.S.C. §1407.

shares, or the parties agree about the percentage claims of the other parties. The question is how to design a division rule to ensure that all parties feel that they have received their agreed-upon share. The attorney fee division context is different because there is no *a priori* agreement on how much each attorney deserves. One might argue, however, that the implicitly agreed-upon share is related to a party's desert, which is in turn related to the party's relative contribution to the litigation. We will examine the limited fair division literature relevant to this context below.

The second goal of the fee division problem is to provide proper incentives among the members of the attorney group to represent their clients' interests zealously. Recall that the mass litigation context features absent clients and a court ill-equipped to monitor attorney behavior. So just as we worried about incentives in the single-attorney context, we again must worry about incentives in the multiple-attorney context. This incentive goal, to be sure, is related to (and perhaps can be simultaneously solved with) the fairness goal, but they are not the same.

The most useful prior work on the incentive problem comes from the industrial organization literature. This literature proposes various compensation schemes to overcome the problem of agency costs in single-principal-multiple-agent contexts and reaches several conclusions relevant here. First, it suggests that if the work product of each member of the team is dependent on the performance of the other members, then team member compensation should depend on the team's overall performance (Baiman & Demski 1980; Holmstrom 1982; Varian 1990). Note that in the main, this situation holds true in MDL. For example, the success of the lawyer handling a summary judgment motion depends on the lawyer handling discovery. Note also that modern court practice follows the literature's prescription. Courts typically use the percentage method to establish the pot of money for attorney fees, which is a measure of the team's overall performance. As long as our ultimate division rule involves splitting the pool further by percentage, we will maintain a link between individual compensation and overall team performance.

Second, the industrial organization literature suggests that team member compensation should depend not only on the team's overall performance, but also on each member's relative contribution to the effort (Itoh 1992; Macho-Stadler & Pérez-Castrillo 1993; Che & Yoo 2001; Siemsen, Balasubramanian, & Roth 2007; Baldenius, Glover, & Xue 2016; Gershkov, Li, & Schweinzer 2016). Otherwise, we would generate a moral hazard problem: if team members are compensated only by how well *others* do, then they will shirk on what they contribute themselves (Alchian & Demsetz 1972; Lazear & Rosen 1981; Mookherjee 1984; Landeo & Spier 2015). The difficulty is how to assess the relative contribution of each team member. Ideally, compensation should be

proportional to a team member's relative contribution. A lawyer that is *twice* as responsible as another for the team's production should be awarded *twice* the compensation. If that is not possible, then imprecise measures of relative contribution like an ordinal ranking of who contributed the most may be just as good or are at least better than nothing (Lazear & Rosen 1981; Green & Stokey 1983; Holmström & Milgrom 1990; Gershkov, Li, & Schweinzer 2016).

How can principals figure out each team member's relative contribution? Most of the literature is unhelpful on this question; it often assumes some "technology" to do it (Gershkov & Winter 2015). The most helpful papers, however, suggest one of three methods: (i) permit team members to decide amongst themselves what their relative contributions were (Itoh 1993)¹³; (ii) give one team member a residual (i.e. what is left over after the others are paid) and ask her to monitor and pay the others (Alchian & Demsetz 1972); and (iii) use multistage techniques to induce team members to make honest reports of each other's contributions (Ma, Moore, & Turnbull 1988; Varian 1990). Underlying many of these tactics in one way or another is the same idea: relying in some way on team members themselves to determine what their relative contributions were; that is, relying on peer assessments (Alchian & Demsetz 1972; Kandel & Lazear 1992; Itoh 1993; Gershkov & Winter 2015; Deb, Li, & Mukherjee 2016). We will return to this idea in the discussion that follows.

In summary, our conceptual analysis of the fee division problem suggests that our solution must appear fair and provide sufficient incentives. To encourage acceptance and to avoid wasteful secondary litigation, the solution should be equitable and sufficiently transparent to avoid suspicion. To create proper incentives, the solution should tie individual compensation to the team's overall performance and the relative contributions made by each team member. Both goals in a sense demand an accurate method for determining relative contribution. It is thus no accident that all the approaches we explore below attack this problem of relative contribution.

3.2 Existing Approaches

Under current practice, courts handle fee division in two principal ways. The first is to let the attorneys split the fees themselves. The second is to use the lodestar or adjusted lodestar method.

13 Some papers find that delegation undermines the ability of the principal to induce optimal performance because the agents can collude against the compensation scheme (Holmström & Milgrom 1990). It is not clear to us if this is possible if the pool of compensation starts from a percentage of the team qua team's output.

Letting the attorneys sort it out on their own is fine as long as it works. Usually, judges appoint one lawyer or a small subcommittee to propose an allocation, and if none of the other team members objects, then that allocation sticks.¹⁴ Such mutual consent, however, is rarely forthcoming. Deferring to a small minority to divide the pie encourages self-dealing, and it is no exaggeration to say that the allocation of fees in MDL has led to fights that are even uglier and more ferocious than the original litigation (Becker, Specter, & Kline 2016; Bronstad 2019a; Bronstad 2019b). This strategy is therefore decidedly suboptimal or at minimum, unreliable. It provides no guarantee against protracted secondary litigation, and the uncertainties surrounding the expected compensation likely harm attorney incentives to do good work.

The lodestar method—using hourly rates and hours worked as a measure of relative contribution—also has serious shortcomings. As previously noted, paying lawyers based on their lodestars gives them incentive to be inefficient with their time. Lodestar is also a poor proxy for relative contribution, as not every hour has the same impact on a litigation.¹⁵ This is precisely why the law-and-economics literature rejected lodestar for calculating single-agent compensation.

Judges have tried to adjust the lodestar method to reflect the heterogeneity of work. For example, the lawyers who review documents during discovery often see their lodestars discounted relative to lawyers who do other things.¹⁶ But to determine the relative value of various tasks, courts have had to resort to questionable methods. Since judges are ill-positioned to monitor the attorneys directly, they frequently ask the lawyers to describe their contributions

14 See *Victor*, 623 F.3d at 85 (providing lead counsel with discretion to allocate the attorneys' fees); *In re Vitamins Antitrust Litig.*, 398 F. Supp. 2d at 213–14 (giving the three lead co-counsel authority to split the attorneys' fees into three pools and allocate to plaintiffs' lawyers accordingly); *In re Initial Public Offering*, 2011 WL 2732563, at *1 (describing the method where one firm acted on behalf of a committee of lead counsel in order to propose a fee allocation); Order and Reasons Allocating Common Benefit Fees in the Knauf Aspect of this Litigation, *In re Chinese Drywall Prods. Liab. Litig.*, No. 2:09-md-02047-EEF-JCW, at 7-8 (E.D. La. February 4, 2019) (describing the method of appointing a Fee Allocation Committee to recommend the appropriate allocation for the team of attorneys).

15 Many papers in the economics and industrial organization literature assume that the ideal would be to base compensation on each agent's "effort." (Macho-Stadler Pérez-Castrillo 1993). But is not the number of hours worked equivalent to effort and therefore would not the lodestar method be a good method according to this literature? The answer is no. The models in this literature assume that a given level of effort necessarily produces a certain outcome: for example, low effort produces a weak outcome and high effort produces a good outcome. But in the real world there is no necessary correlation between how many hours people work and what kind of outcome they achieve.

16 See, for example, *In re TFT-LCD Litig.*, 2013 WL 1365900 at *9 ("Factors meriting an upward adjustment include. . . performing higher-skill tasks (e.g., taking depositions). . . Factors meriting a downward adjustment include. . . performing largely document review. . .").

after the litigation has concluded.¹⁷ But unless team members expect to be on the same team in the future (a rare occurrence¹⁸) (Landeo & Spier 2015, p. 509), they have every incentive to inflate their own contributions and denigrate those of others. Alternatively, courts make adjustments based on the opinion of the lead lawyer or the steering committee. These similarly self-serving reports, however, are a poor source of information and generate resentment.

3.3 Manufacturing a Client

As the reader will recall, the fundamental problem with attorney fees in the mass litigation context is the absence of an involved, sophisticated client. One solution is therefore to simply manufacture one, and that is precisely what Charlie Silver and Geoff Miller propose. They propose that the court appoint the lawyer with the most clients in the MDL to be the lead lawyer. That lawyer cannot do any work for the common benefit of the MDL (to prevent self-dealing), but that lawyer decides who will do common benefit work and how much those lawyers will eventually be paid for the common benefit work. The lead lawyer then taxes all the cases—including his own—at the same rate to pay for the common benefit work so retained (Silver & Miller 2010).

Under the Silver–Miller proposal, the lead lawyer has strong incentives to pick, compensate, and monitor attorneys judiciously. He will want to hire good lawyers for the common benefit work and monitor them because he has a lot of cases in the MDL and wants them to succeed because he will collect a contingent fee on each of them, usually between 33 and 40 percent. The lead lawyer also wants to hire good lawyers as cheaply as possible because, as the lawyer with the most cases, the lead lawyer will pay more tax money than anyone else (Silver & Miller 2010). This proposal thus solves the problem of how much money should be paid to the other common benefit lawyers (i.e. the size of the “pool”) and how it should be allocated. Harkening back to the industrial organization literature, this proposal basically pays a residual to a monitor; in this case, the residual is the contingency fees the lead lawyer will earn on all his own cases minus the tax he has assessed on himself and all the other lawyers.

The Silver–Miller solution is innovative, but it has several drawbacks. First, in order for the strategy to work, the appointed monitor cannot assign the

17 See *id.* at *7–16 (allocating the attorneys’ fees with assistance by a Special Master). For examples where the court allocated the funds after objections are made by some of the lawyers, see *Victor*, 623 F.3d at 84, *In re Vitamins Antitrust Litig.*, 2011 WL 2732563 at *214, and Order and Reasons, *supra* note 14, at 8–9.

18 It is rare that the team of lawyers in one MDL is the same as a team in another MDL, but it is more common for *some* members of a team to reappear in other cases (Burch 2019, pp. 235–237). We are unaware of any models that incorporate this sort of partial repeat playing into their parameters.

pretrial work to itself. This means the appointed firm necessarily loses a significant stream of potential work and revenue. There is no guarantee that the residual will be as lucrative as doing the actual legal work itself. Second, in many consolidated MDLs, no single law firm has an overwhelming portion of clients, so those strong monitoring incentives may be difficult to find. Third, the proposal is an *ex ante* schema. A court must have the foresight to appoint the monitor ahead of time. Silver–Miller can only prevent attorney fee disputes in the future, it does not help resolve attorney fee disputes after they arise.

Finally, whatever its merits, courts have been reluctant to adopt Silver–Miller in practice because, as we understand, the firm with the largest stake is often an “advertising lawyer,” whom courts view disparagingly as being more skilled at recruiting clients than litigating cases. In our view, this objection is inapposite. The “client” that Silver–Miller attempt to manufacture need not be a top-flight litigator to be a good monitor. Indeed, in ordinary litigation, the clients who perform this monitoring function are often not lawyers at all. Nevertheless, this apparent judicial preference for selecting lead attorneys who do the best legal work remains an obstacle to Silver–Miller’s acceptance.

4. A PEER ASSESSMENT APPROACH

In our view, the solution to the fee division problem does not lie in any of the previously explored methods. It does not rest in using an adjusted lodestar or manufacturing an artificial client. Instead, we propose to follow the insights from the industrial organization literature and use peer assessment. The best source of information on a firm’s performance is often the other team members. Using team member reports also imbues the process with a participatory or democratic character that promotes eventual acceptance. Indeed, taking the idea one step further, the methods we ultimately propose are flexible enough (if so desired) to incorporate assessments not only of peers, but also self-assessments and assessments by the court.

4.1 Prior Work

4.1.1 Brams–Taylor Game Theoretic Approach

One of the earliest methods to divide fees using peer assessments is a clever algorithm proposed by Steven Brams and Alan Taylor in the game theory literature. [Brams & Taylor \(1996\)](#) propose the following algorithm for how a team of workers should divide an unexpected bonus:

- (1) Ask each team member to propose an allocation for the team’s bonus.

- (2) Derive a baseline allocation for each member by averaging the allocations he is awarded by the other members.
- (3) Pay each member exactly what he proposed for himself, but in order of least “greedy” to most “greedy,” with greedy defined as the difference between the self-allocation and the baseline calculated in (2). Allocations are paid until the pool runs out (Brams & Taylor 1996).

Overly greedy members risk a shortfall, since they are paid last, and the pool may run out of funds. The game therefore drives each team member toward allocations that match what they think *others* will want to give them, which are assumed to be the most accurate reports.

Brams and Taylor do not give a full game theoretic analysis of this game, but some aspects of the game are clear. First, if all the parties agree on an allocation, then there will be an equilibrium¹⁹ in which each firm proposes the generally accepted allocation for himself. Thus, there are an infinite number of equilibria. Since there is no objective true allocation it is hard to formally analyze the optimal strategies for the firms. Indeed, it is hard to argue that there is even a natural focal point for the firms to coalesce around. If there were a natural focal point the division of attorney fees would not be so fraught in the first place!

On the other hand, intuitively, Brams–Taylor creates strong incentives to make self-reports that are close to the consensus of the other firms, and participants in a sense cannot complain about the results—either they receive what they request, or they are punished for their greed. We think, however, that Brams–Taylor suffers from several features that limit its usefulness in MDL. One problem is that it potentially leaves one or more lawyers with a massive shortfall (or indeed nothing!) after years of hard work. This nightmare scenario violates notions of fairness and desert. The affected party’s greed or overly high opinion of itself may be partly to blame for its own shortfall, but the punishment is completely out of proportion with the offense. Perhaps this scenario will never happen because fearful team members will always request too little rather than too much, but what if it does happen? Will a judge really give that team member nothing? If the judge sympathetically deviates from the harsh result, then the judge destroys the very discipline that Brams–Taylor needs to make it successful in the first place.

The harshness of the null result also likely makes Brams–Taylor unacceptable to the attorneys themselves. Who would willingly agree to run this risk to their livelihood after years of work? Would not attorneys come away resentful for having been forced to deliberately undervalue their own work to avoid the nightmare scenario? And might these risks harm the incentives attorneys have

19 Technically a Nash equilibrium in pure strategies.

to work hard? The brilliance of Brams–Taylor is that it converts a very difficult problem into a game. The downfall of Brams–Taylor is that given the money at stake, fee division is anything but.²⁰

The second problem with Brams–Taylor is that the method assumes that all team members know enough about each other to rate each other at all. Some MDL teams can consist of dozens of lawyers and not everyone will know what everyone else has done. So, while Brams–Taylor could provide a solution in small cases involving smaller stakes, we doubt that it can truly solve the fee division problem in large-scale MDLs.

4.1.2 de Clippel, Moulin, and Tideman’s Algorithm

Another approach is taken by Geoffroy de Clippel, Herve Moulin, and Nicolaus Tideman (hereinafter “CMT”). They propose a method of fair division that relies on each team member rating the relative contributions of each pair of *other* team members (de Clippel, Moulin, & Tideman 2008). Suppose there is a team of three members: Firms A, B, and C. Firm A would be asked what the contribution of B was relative to C; B would be asked what the contribution of A was relative to C; and C would be asked what the contribution of A was relative to B. CMT derive a family of formulas that take these pairwise ratings and produce an overall allocation with strong mathematical properties.²¹ For example, under their schema, no firm is able to influence what it ultimately receives,²² and if there is an allocation that is consistent with all of the reports, their formula will find it.

To be more concrete, suppose Firm A reports that B’s work was worth twice C’s. Firm B reports that A’s work was worth three times C, and C reports that A’s work was worth 1.5 times B. There is a division that is consistent with all three of these reports, and CMT finds it: A gets 1/2, B gets 1/3, and C gets 1/6.

Because a firm’s allocation is based solely on the reports of the other firms it follows that all sets of reports form an equilibrium. So, as a game, there is little to analyze in CMT. It is worth noting that, as in the situation of Brams–Taylor,

20 We should note Brams–Taylor can also suffer from the opposite problem. If the firms collectively underestimate their shares (perhaps to avoid the nightmare scenario), part of the attorney fees pool will remain unallocated. Such an outcome is not an equilibrium in the game-theoretic sense but it may arise in the real world. Even if it does arise, however, the court could simply return the unallocated fees back to the plaintiffs.

21 For three firms, there is a unique formula: if we let r_{BC} be the report from A about the relative contribution of B to C, and similarly r_{AC} is B’s report of the relative contribution of A to C, then de Clippel’s allocation to C is $\frac{1}{1+r_{BC}+r_{AC}}$. Note that this allocation is not dependent on any report from C. Similar formulas give the allocations to A and B. For four or more firms, de Clippel derives a family of solutions, but they are complex and omitted here for brevity.

22 For an expository presentation of CMT’s formula and theory, see Tideman & Plassmann (2008).

if there is an allocation that is consistent with all of the reports, then that is the allocation produced by CMT as well.

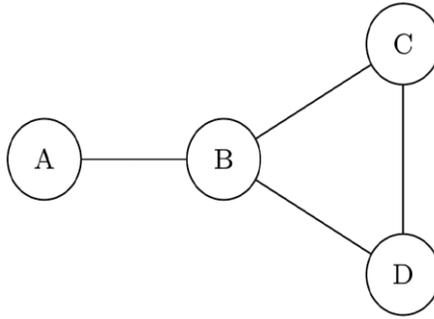
Unfortunately, CMT has some attributes that make it unattractive for solving the fee division problem. One significant weakness is that the formulas are rather difficult to explain or motivate. To be readily accepted by the attorneys, one has to be able to persuade them why a rule chooses the division that it does, and doing so with CMT is challenging. One can of course appeal to CMT's mathematical properties, but the intuition behind the formulas is obscure. When there is no allocation consistent with all of the reports (the most likely scenario), what tradeoffs does CMT make among the disagreeing reports? Consider a slight modification to the example above: Firm A reports that B's work was worth twice C's, B reports that A was worth three times C, and C reports that A was worth *the same* as B. There is no division that is consistent with all three of these reports, but yet the CMT formula produces an allocation: $A = 3/7$, $B = 2/5$, and $C = 1/6$ (de Clippel, Moulin, & Tideman 2008, Prop. 1).²³ Why? We do not really know, and that opacity leaves us (and we suspect will leave courts and attorneys) uneasy. Firms will want some explanation for why the outcomes are different from what they recommended, and CMT fails to provide it.

The other significant problem with CMT is that like Brams–Taylor, CMT deals poorly with missing information. For CMT to work, every possible pair of team members must receive at least two relative-contribution reports.²⁴ These significant informational demands may be difficult to meet in MDL, where incomplete data will be common. Consider the network in Figure 1. If Firms B, C, and D work together, but Firm A only works with Firm B, CMT will not work. This kind of situation will become frequent as the number of firms grows large.

Like CMT, our proposed approach to the fee division problem is also based on peer assessments, but we eliminate the concerns raised by the CMT method by attacking the problem differently. We propose calculating the final

23 The astute observer will note that the sum of the shares allocated to the firms do not add to 1. That is, the de Clippel rule is not efficient in the sense that it forces some of the good to go unallocated. This would seem problematic. One could of course normalize the allocations, but the normalized allocation would no longer be guaranteed to have CMT's strong mathematical properties. As it turns out, however, if there are four or more firms, a generalization of the de Clippel rule will in fact allocate all of the good (de Clippel, Moulin, & Tideman 2008, Thm. 2).

24 CMT can work with one relative-contribution report for each pair, but then it may not fully allocate all of the money. For CMT to be "efficient" and divide the entire pool, two reports are needed for every possible pair.

Figure 1. Example firm network.

allocation in two related, but different ways. One is to treat the calculation as a pure optimization problem: what allocation “fits” (in the sense of minimizing squared error) the various peer assessments best? The other is to treat the calculation as a Bayesian modeling problem: if we assume that the peer assessments are observations (with error) of some underlying true set of relative contributions, what is that underlying distribution most likely to be?²⁵

4.2 Problem Specification

Suppose that N firms, labeled $1, 2, \dots, N$, have worked together on a litigation matter. A court has issued a collective award of attorney fees, and the objective is to distribute the sum according to “desert” as defined by the firms. To begin, we ask each firm to rate the relative contribution made by each of the other firms. We label these observed ratings as S_{ij} , $0 \leq S_{ij} \leq 1$, where i is the rater and j is the rated firm. To prevent self-dealing (for now), a firm may not rate itself. In addition, because some firms may not have sufficient contact with some of the other firms to make an educated rating, some of the S_{ij} may be missing. So, for example, we might have a score matrix that looks like [Table 1](#).

Note that under this construction, the row sums ($\sum_j S_{ij}$) necessarily equal 1, since each rating firm makes *relative* assessments among the firms for which it has information. Missing values are NA, as distinct from a zero contribution. Given this dataset, the goal of the approaches below is to arrive at a justified

25 Those with some theoretical familiarity with regression may recognize a conceptual analogy. One common method of estimating linear regressions is “least squares,” in which we seek the regression line that minimizes the sum of square errors between the fitted line and the data. That is an optimization problem. Another view of linear regression is to treat the data as observations with error from a true model. Under this view, we estimate the coefficients (again, the fitted line) that are most likely—hence, the term “maximum likelihood estimate.”

Table 1. Example score matrix

Rater/rating	Firm 1	Firm 2	Firm 3	Firm 4
Firm 1		0.50	0.40	0.10
Firm 2	0.70		0.30	NA
Firm 3	0.85	NA		0.15
Firm 4	NA	0.50	0.50	

estimate for the contribution made by each firm to the litigation as a whole. We will denote by α_j the final recommended allocation for firm j .²⁶

4.3 Optimization

Our first approach to the fee allocation problem is optimization. In this approach, we try to find an allocation that best “fits” all of the reports by the firms. To accomplish this, we choose allocations $\{\alpha_j\}$ so as to minimize the error between our allocations and the reports provided by the firms themselves. We will analyze two different measures of error. The first is based on pair-wise error and the second is based on the error of each individual firm assessment. These measures will be explained below. To be clear, in this approach we do not assume that there is any particular “true” allocation. We merely try to find an allocation that best fits with all of the reports of the firms.

4.3.1 Pairwise Error

Suppose that firm i reports on the relative contributions S_{ij} and S_{ik} of the two firms j and k . Then we would want our final allocations α_j and α_k to satisfy:

$$\frac{S_{ij}}{S_{ik}} \approx \frac{\alpha_j}{\alpha_k}.$$

A natural measure of the error produced by our allocation relative to firm i 's assessment of firms j and k is as follows:

$$\left(\frac{S_{ij}}{S_{ik}} - \frac{\alpha_j}{\alpha_k} \right)^2$$

but this measure suffers from a number of drawbacks. First, this measure is undefined if some of the ratings $S_{ij} = 0$. Even if we rule out 0 ratings, if some of the ratings are very small then the resulting optimization problem (described

26 Further note that even a simple example like Figure 1 already violates the de Clippel informational requirements. For example, Firms 1 and 2 are never rated together.

below) can be very unstable, that is, there can be wild swings in the allocation for small changes in the ratings. Second, the measure is not symmetric in the firm labels j and k .²⁷

Alternatively, we could choose the symmetric measure

$$(\alpha_k S_{ij} - \alpha_j S_{ik})^2.$$

This formulation has a number of advantages. It is symmetric in the labels, allows for 0 ratings, and it leads to a well-behaved quadratic optimization function. Algebraically we have

$$(\alpha_k S_{ij} - \alpha_j S_{ik})^2 = (\alpha_k S_{ik})^2 \left(\frac{S_{ij}}{S_{ik}} - \frac{\alpha_j}{\alpha_k} \right)^2.$$

So, this latter measure amounts to a weighted version of our nonsymmetric measure. This weighting will magnify the error when both the report S_{ik} and the allocated amount α_k are large. The former may happen because the reporting firm only has information on a few other firms, thus the relative values will tend to be larger, or because this particular firm merited a large allocation. In either case, we think it is a plus that large reports are held to stricter requirements and so forcing the error they produce to be smaller.

Summing over all firms and their reports we are left with the following quadratic optimization problem²⁸:

$$\begin{aligned} &\text{Minimize } \sum_{i=1}^N \sum_{\substack{j,k \neq i \\ S_{ij}, S_{ik} \text{ defined}}} (\alpha_k S_{ij} - \alpha_j S_{ik})^2 \\ &\text{subject to } \sum_i \alpha_i = 1 \\ &\alpha_i \geq 0 \forall i \in \{1, 2, \dots, N\} \end{aligned}$$

4.3.2 Individual Error

Another approach would be to focus on the individual assessments of the firms. Suppose firm i reports a relative share S_{ij} for firm j . Then we would want

27 A referee has pointed out that we could correct the symmetry issue by using the measure $\left(\frac{S_{ij}}{S_{ij}+S_{ik}} - \frac{\alpha_j}{\alpha_j+\alpha_k} \right)^2$. However, this does not fix the problem of reports of 0.

28 This is a classic, well-behaved quadratic optimization problem which can be solved using only linear algebra (Avriel 1976, p. 186). Moreover, the nonnegativity conditions on the α_i are unnecessary, since one can show that the unique solution will always satisfy those conditions.

$$S_{ij} \approx \frac{\alpha_j}{\sum_{\{k|S_{ik} \text{ defined}\}} \alpha_k}.$$

Reasoning as before we might wish to quantify the error in this allocation as follows:

$$\left(S_{ij} - \frac{\alpha_j}{\sum_{\{k|S_{ik} \text{ defined}\}} \alpha_k} \right)^2.$$

Unlike the earlier measure, this one exhibits no particular asymmetry. It is, however, incompatible with allowing reports of 0 as well as being rather ill-behaved for optimization purposes. It also has the property of weighting all errors equally, even if the actual allocations involved are rather small. We can create both an easier optimization problem as well as a more meaningful weighting by using the following equation:

$$\left(S_{ij} \sum_{\{k|S_{ik} \text{ defined}\}} \alpha_k - \alpha_j \right)^2 = \left(\sum_{\{k|S_{ik} \text{ defined}\}} \alpha_k \right)^2 \left(S_{ij} - \frac{\alpha_j}{\sum_{\{k|S_{ik} \text{ defined}\}} \alpha_k} \right)^2$$

as our quantification of the error in the allocation for firm *i*'s estimate for the relative share for firm *j*. Note that now we are weighting the error by the amount of the allocation that firm *i* observes. This leads us to the well-behaved quadratic optimization problem.²⁹

$$\begin{aligned} &\text{Minimize } \sum_{i=1}^N \sum_{\substack{j \neq i \\ S_{ij} \text{ defined}}} \left(S_{ij} \sum_{\{k|S_{ik} \text{ defined}\}} \alpha_k - \alpha_j \right)^2 \\ &\text{subject to } \sum_i \alpha_i = 1 \\ &\alpha_i \geq 0 \forall i \in \{1, 2, \dots, N\} \end{aligned}$$

4.4 Bayesian Model

An alternative approach to the problem is to use a statistical model in which each firm's "true" contribution to the litigation is modeled as a latent variable, α_j . Recall that contributions are measured on a relative basis, so the vector α is a unit simplex, that is, $\alpha_j \geq 0$, and $\sum_j \alpha_j = 1$. The α_j 's are latent and therefore not directly observed. Instead, we observe the relative ratings given by other firms, which provide information about the α_j 's indirectly and with error.

29 Like the previous optimization problem, this problem is well-behaved and can be solved using only linear algebra. It is also still true that nonnegativity condition is not required as it will always be satisfied at the unique optimum.

In what follows, we discuss Bayesian models of increasing sophistication to estimate α_j .

4.4.1 Linear Approach

One straightforward approach is to view the observed scores (S_{ij}) as consisting of the “true” contribution of firm j relative to all of the other firms rated by rater i plus some Gaussian error. So, for example, if we let i be the rater firm, j be the rated firm, and X_{ik} be an indicator variable for when firm i has sufficient contact with firm k to evaluate its contribution, then the observed score can be modeled as follows:

$$S_{ij} = \mu_{ij} + \epsilon_{ij}$$

where $\mu_{ij} = \frac{\alpha_j}{\sum_{k \neq i} \alpha_k X_{ik}}$ and $\epsilon_{ij} \sim N(0, \sigma_j^2)$.³⁰

With this model, each firm’s contribution can then be estimated using Markov Chain Monte Carlo (MCMC) methods that are standard to Bayesian analysis. The scores given by the firms constitute the observed data (S_{ij}), and the underlying firm contributions are the unknown parameters of interest (α_j). The MCMC methods ascertain what distributions of parameter values are most likely to have given rise to the observed data. We then use the means of these distributions as the estimates for the parameters. In our study, estimation was done using R in conjunction with the STAN Bayesian modeling platform (Stan Development Team 2020).

4.4.2 Compositional Approach

A significant limitation to the linear approach is that it fails to account for correlation among the observation errors. Because α is a simplex, error in measuring one firm’s contribution should negatively correlate with the error in measuring other firms’ contributions. For example, suppose Firm A rates the relative contributions of Firms B and C. If A underestimates B’s contribution, then it must necessarily overestimate C’s contribution, since the total relative contributions of B and C must sum to 1.

The log-ratio approach to compositional data proposed by Aitchison (1986) helps model this correlated error. Let’s assume that we have N firms involved in the division, and that the true (latent) contribution for each firm is represented by α_j , where $j = 1, \dots, N$. Because the vector α is a simplex, it is

30 Here we give the α_j ’s a flat Dirichlet prior, since they constitute a simplex, and σ_j ’s receive uninformative inverse-gamma priors (Gelman 2006).

completely defined by its first $n = N - 1$ elements. So instead of focusing on the individual α_j 's, we can focus on their log-ratios, namely:

$$\mu_j = \log\left(\frac{\alpha_j}{\alpha_N}\right), \quad j = 1, \dots, n.$$

We can then view the observed log-ratios to be the true log-ratios plus some error term. Let s_{ij} be the observed contribution of firm j as judged by firm i , and the vector \mathbf{s}_i contain all of the contributions observed by firm i except the last term (s_{iN}), namely, $(s_{ij}, j = 1, \dots, n)$. Then, we can model those contributions as follows:

$$\log\left(\frac{\mathbf{s}_i}{s_{iN}}\right) = \boldsymbol{\mu} + \boldsymbol{\epsilon}_i,$$

where i indexes the firm doing the judging, and the vector $\boldsymbol{\epsilon}_i \sim \mathcal{N}^n(0, \boldsymbol{\Sigma})$ models the correlated error.

One can then extend this compositional method to account for missing data, such as when one firm lacks sufficient information to rate another firm's contribution. Additional technical details can be found in Appendix A. Once again, the unknown parameters (α_j) can be estimated using MCMC techniques.

4.4.3 Random Effects

A limitation of the earlier Bayesian models is that they treat the rater firms as interchangeable scientific instruments, so that each set of observed ratings is like any other. However, rater firms are bound to have idiosyncrasies. Certain raters may be more knowledgeable or competent, and thus they may have lower error when estimating contributions.

Of even greater concern, several firms may attempt to collude with each other—for example, they may agree to rate each other at levels far in excess of their desert. To be sure, we do not know how often collusion will occur. For one thing, if peer reports are kept confidential (as they likely will be), colluders will be unable to verify their counterparty's behavior, a flaw that usually undermines such schemes. For another, if such collusion was discovered, the reaction and resulting sanctions from courts are likely to be severe.

One way to address these concerns is to use a random-effects model.³¹ Rather than treat each firm as interchangeable, a random effect acknowledges

31 A less technical option is for the court to randomly drop some of the peer assessment data as a matter of course. This practice would make the benefits of collusion less of a sure thing, reducing the incentives to engage in it. Since our methods tolerate missing data well, the anti-collusive benefits of this random-drop strategy may outweigh the corresponding loss in data.

that each rater firm has its own variability or reliability. The rater firms, however, all come from a common distribution. So, for example, most firms may exhibit some average level of error, but a few may be especially knowledgeable and accurate, while a few others may be unusually erratic.

By modeling each firm separately, a random-effects model directly accounts for accuracy differences among the different raters. A random-effects model also helps resist collusion because (to be effective) ratings given by colluders will depart strongly from the consensus. The model will interpret such departures as errors, causing the model to view the colluder as less reliable and to downweight its ratings.

Consider this example: suppose that there are ten firms in the pool, and Firms 4 and 5 collude to give higher ratings to each other. Since Firm 4's estimate of Firm 5's contribution will deviate from the assessments made by the other eight firms, the model will estimate Firm 4's variability to be quite high, effectively downweighting Firm 4's observations. A similar downweighting will occur due to Firm 5's unorthodox estimate of Firm 4's contributions. Such collusion protection is less effective when the number of firms is small, because detecting "outlier" ratings will be difficult if not impossible. But a random effects model provides some incentive and assurance against collusive behavior when there are a substantial number of firms.

Collusion is also possible among more than two firms, but the principle of the random-effects model would be the same. Firms departing from the consensus will find their ratings downweighted. Naturally, the method will have trouble thwarting large-scale collusion, since it then becomes difficult to distinguish colluders from good-faith participants. Large groups of colluders, however, presumably experience a higher probability of leaks or defections, and can be punished externally.

Technical details about implementing random effects in the compositional model can be found in Appendix A.

5. DISCUSSION

5.1 Generally

Both the optimization and statistical methods offer promising solutions to the fee division problem. To start, the methods are based on peer assessment, making them more participatory and democratic. The methods are also straightforward to motivate and explain. Least-squares optimization and Bayesian modeling are well-known statistical techniques. The optimization method simply chooses an allocation that best fits (i.e. minimizes deviations from) the

reports of the firms.³² The Bayesian model finds the most probable “true” allocation given the variation in the law firms’ reports. Unlike Brams–Taylor, our methods never impose massive shortfalls on team members for poorly estimating their own contribution. A team member may receive a smaller than expected share, but only because everyone else says it should. The proposed methods are also capable to handling incomplete data and are robust to collusion.

In simulations, which are detailed in Appendix B, our methods consistently yield allocations that accord with each other and that are close to the “true” allocations that gave rise to the simulated data.³³ This agreement occurs whether one uses an optimization or Bayesian approach, whether the data are complete or incomplete, and whether some of the actors collude or not. Such consistency should provide courts and attorneys with confidence that our methods’ suggestions are stable and fair estimates of the team members’ contributions. To the extent that our various methods differ in their conclusions, courts can either average them together, or decide *ex ante* to use one of them exclusively. Courts could also choose to use them as guideposts (or justifications) for a more subjective determination.

5.2 Strategic Behavior and Self-Reporting

One advantage that the CMT method theoretically has over our methods is that it is impervious to strategic behavior by raters. De Clippel, Moulin, and Tideman mathematically prove that under their method, a firm is entirely unable to affect its own allocation through its ratings of other firms.³⁴ While strategic rating is theoretically possible under our proposed methods, it is practically unlikely. How a firm’s reports about *other* team members affects its own share is extremely difficult to predict, particularly because it depends on how the other firms are going to rate each other as well. Additionally, the dense “web” of relative assessments prevents any one set of scores from affecting the

32 Unlike the Bayesian method, the optimization method requires no assumptions about an underlying objective truth. Philosophically, one can view the optimization method as the ultimate “split-the-difference” solution—finding the best midpoint among all of the ratings.

33 The close agreement between the optimization and Bayesian results may initially seem remarkable, but it should perhaps be unsurprising. As we know from the Gauss–Markov Theorem in the linear regression context, the least squares estimator (optimization approach) is the best linear unbiased estimator (statistical approach). Perhaps what we have here is a compositional data variant of the Gauss–Markov Theorem, although further research would be needed to ascertain which precise methods are linked and under what conditions.

34 Indeed, CMT mathematically prove that their method is the only method that can solve the fair division problem while satisfying all three axioms that they set out: objectivity, strategy-proofness (the one we are discussing), and consistency.

final allocation by much. If collusion barely works (as seen in our simulation results), strategically scoring others in an attempt to indirectly benefit oneself is nearly impossible.

Being completely impervious to strategic behavior may also come at too high a cost, because it necessarily means prohibiting self-reports. There is a participatory value in allowing firms to provide at least some direct input about their own contribution to the litigation. It enables firms to exercise a right to be heard and promotes acceptance of the process. In many ways, this is the reason why the “I cut, you choose” cake-cutting algorithm, or even the Brams–Taylor method, has intuitive appeal.

Prohibiting self-reporting also carries informational costs. The firm with arguably the best information about what Firm A did during a litigation is Firm A itself. Banning self-reports may therefore throw away valuable information. More importantly, if firms are unable to self-report, what incentive will they have to make careful assessments about other firms? When a firm engages in self-reporting, it is highly motivated to get the other firm contributions right, because its own contribution will be assessed relative to those contributions. When a firm cannot self-report, its incentives to assess other firms with precision becomes more tenuous.

The good news is that our methods are capable of handling self-reporting.³⁵ It handles self-reports like all other reports. And just as our methods exhibit natural resistance to collusion, they can also resist distortions caused by overly generous self-reports. Using the collusion resistance model provides still further protection, because it penalizes self-reports that significantly depart from the community consensus. If a firm wishes its own report to be taken seriously, it would be wise to rate itself and the other firms carefully and accurately.³⁶

Simulations of our methods in the presence of self-reports bear out these predictions. The models are unfazed by the addition of self-reports, even if unabashedly self-interested.

5.3 Data Concerns

A fee division method is only as good as its data, and so we finish our discussion with some caveats about data. Although our proposal has significantly more relaxed data requirements than previous proposals, it still requires some degree of interconnection between the various ratings. It is unlikely that any

35 Our thanks to Ben Alarie for pointing out this possibility.

36 A referee noted another appealing prospect of allowing self-reports. We could use our method (without self-reports) to get a base-line allocation and then use the self-reports to penalize firms that over-claim relative to this baseline, à la Brams–Taylor.

firm will have no ratings, but small-share firms may have only one point of connection to the broader group—for example, recall [Figure 1](#), in which Firms B, C, and D all work together, but A only interacts with B. Such small-share firms are more vulnerable to the whims of their one rater and do not have an entire network protecting them. In many ways though, this problem is not unique to our method. When no one really has information about Firm A’s contribution, Firm A is left vulnerable under almost any fee division scheme.

Another structure that could create problems is a dumbbell-shaped network where one firm forms the link between two larger groups of firms. In this situation, the relative weights within the groups may be well estimated, but the weight between the two groups will depend entirely on the linking firm’s opinions. Whether this dumbbell network ultimately creates acceptance problems though is unclear. After all, we suspect that a firm’s satisfaction with its allocation depends on how its portion compares to firms with which it is familiar, not portions received in the foreign half of the network.

Peer reports can also be afflicted by various forms of bias. For example, a firm’s prominence may inflate other firm’s assessments of its contribution. This phenomenon may in turn cause firms to preen, wasting resources. More perniciously, racial or gender bias may cause undervaluation of minority or female lawyers. And on a more technical level, it may be more cognitively difficult to make accurate distinctions among smaller contributions (2 vs. 4 percent) than larger ones (20 vs. 40 percent).³⁷ All of these problems, however, are common to all solutions of the fee division problem. They are present whether one applies Brams–Taylor, CMT, or the modified lodestar method (primarily through the multipliers),³⁸ and they are certainly present if a lead attorney or court simply decides on an allocation. To the extent they are ubiquitous, the key is for participants to be mindful of these problems, both in making their ratings and in assessing the answers that the methods provide.³⁹

Finally, we note that nothing in our proposal requires that the raters exclusively be the rated firms themselves. So, if a court chooses, it can add its own assessment of the firms or the assessments of consultants or special masters. For reasons stated in the Part 1, we doubt that courts or special masters will

37 One way to address this problem would be to limit the number of firms rated by each rater. Since raters report relative contributions out of 100 percent, that strategy would reduce the appearance of small values. Note also that the measure of error suggested in Part 4.3.1 punishes errors in large allocations more severely than errors in small ones.

38 Even the plain lodestar method itself fails to avoid potential bias issues. Working excessive hours is analogous to the preening problem, and the use of a firm’s “normal” hourly rate resurrects possible racial and gender disparity concerns.

39 At least for our statistical method, in cases with a large number of firms, it may also be possible to use covariates to check for possible race or gender effects both for raters and ratees.

often have such superior information so as to justify this strategy. Also, a single set of ratings is unlikely to influence the outcome, for the same reasons that make our methods collusion resistant. Nonetheless, courts can use this option as a safety valve, especially if they adjust the model to weight “neutral” ratings more heavily.

6. CONCLUSION

The problem of distributing fees among attorneys in MDL is an important but neglected problem in the legal literature. The existing single-principal-single-agent literature offers a partial solution, but it only tells courts how to determine the aggregate fees owed to the attorneys, not how to divide them. At the same time, the current court practice of relying on a modified lodestar method to aid the division of fees resurrects old complaints about lodestar and its inefficiencies.

In this article, we have surveyed a number of possible innovative solutions to the fee division problem, ranging from Silver–Miller’s appointed monitor, to Brams–Taylor’s game, to CMT’s method of peer assessment. In the end, we believe all of the previous solutions have serious drawbacks that will prevent their successful adoption. Instead, our proposed methods offer a more viable alternative. Our optimization or Bayesian methods are more intuitive to understand, have lower informational requirements, and are resistant to collusion. They should also be easier for the attorneys to accept, as they avoid extreme outcomes and offer the possibility of self-reports. Our proposed methods thus fill an important gap in the literature on how judges can resolve fee disputes at the conclusion of MDL.

Finally, although our focus has been on attorney fee division, the techniques we have reviewed and proposed are not limited to that context. Indeed, they apply to any instance in which team members jointly create value and then subsequently need to divide up the pool. When law firms allocate profits at the end of each year, business associations unwind, or inventors or artists jointly develop creative works, our methods can divide the spoils. Ideally, *ex ante* contracts will govern these divisions, but often they will not. In those cases, our methods can help these actors reach a suitable resolution, whether by themselves or via the courts.

APPENDIX A: COMPOSITIONAL BAYESIAN APPROACH

This appendix develops the technical details for the compositional Bayesian approach, both with and without random effects.

A.1 Basic Compositional Model

As previously suggested in Part 4.4.2, since the contributions for all firms form a simplex, the contribution vector α is completely defined by its first $n = N - 1$ elements.⁴⁰ So instead of focusing on the individual α_j 's, we focus on their log-ratios, namely:

$$\mu_j = \log\left(\frac{\alpha_j}{\alpha_N}\right), \quad j = 1, \dots, n.$$

We can then view the observed log-ratios to be the true log-ratios plus some error term. Let s_{ij} be the observed contribution of firm j as judged by firm i , and let the vector \mathbf{s}_i contain all of the contributions observed by firm i except the last term (s_{iN}), namely ($s_{ij}, j = 1, \dots, n$). Then, we can model those contributions as follows:

$$\log\left(\frac{\mathbf{s}_i}{s_{iN}}\right) = \boldsymbol{\mu} + \boldsymbol{\epsilon}_i, \quad (1)$$

where $\boldsymbol{\epsilon}_i \sim \mathcal{N}^n(0, \boldsymbol{\Sigma})$, and i indexes the firm doing the judging. Here, $\alpha_j, j = 1, \dots, n$, can have uniform priors on $[0, 1]$, or perhaps a weakly informative normal prior (with appropriate constraints) since we know that the relative contributions will tend to be in the lower part of that interval.

For ease of notation, we can borrow the notation from [Aitchison \(1986\)](#) and express the model as $\mathbf{s}_i \sim \mathcal{L}^n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The covariance matrix ($\boldsymbol{\Sigma}$) captures interdependencies among the firm contributions. For example, if Firms A and B worked on the same aspect of a case, then we should expect negative covariance between the error associated with A's contribution and B's contribution.

A.2 Incomplete Data

Thus far, the proposed compositional method only handles complete data. Recall, however, that some firms may not have sufficient information to rate all of the other firms, or we may not want firms to be able to rate themselves. In these cases, the contributions provided by the raters are only the *relative* contributions among the firms for which the rater has enough information.

⁴⁰ $\alpha_N = 1 - \alpha_1 - \dots - \alpha_n$. As [Aitchison \(1986\)](#) shows, the results are invariant to which α_j is chosen to be the α_N —in other words, the results are invariant to permutations of the firms.

Because the log-ratio approach is based on a multivariate normal model, it can handle relative contributions (known as subcompositions) with ease through a linear transformation, as in [Aitchison \(1986\)](#). If $\mathbf{s}_i \sim \mathcal{L}^n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is a complete composition as seen in [Equation 1](#), then subcomposition $\tilde{\mathbf{s}}_i$ has the following distribution:

$$\begin{aligned}\tilde{\mathbf{s}}_i &\sim \mathcal{L}^n(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{\Sigma}}) \\ \tilde{\boldsymbol{\mu}} &= Q\boldsymbol{\mu} \\ \tilde{\boldsymbol{\Sigma}} &= Q\boldsymbol{\Sigma}Q^T,\end{aligned}$$

To where we construct matrix Q based on which indices are chosen for the subcomposition. In particular, let M be the number of firms in the subcomposition with $m = M - 1$, and N be the number of firms in the complete composition with $n = N - 1$. Then

$$Q = F_M Z F_N^T H^{-1},$$

where F_k is the identity matrix I_k with an appended last column of -1 's (i.e. $[I_{k-1} : -j_k]$), Z is the $M \times N$ selection matrix that creates the subcomposition, and H is defined as $H = I + J$, where I is the identity matrix and J is a matrix of ones, as in [Aitchison \(1986\)](#).

A.3 Random Effects

To introduce a random effect into the Bayesian composition model to account for rater variability, we can use a decomposition of the covariance matrix:

$$\boldsymbol{\Sigma} = \boldsymbol{\Omega} D \boldsymbol{\Omega}^T,$$

where D is a diagonal matrix of eigenvalues, σ_j , corresponding to the “scale” of the covariance. We then assume that D has taken the form:

$$D = \gamma_i \tilde{D},$$

where \tilde{D} is the diagonal matrix with elements $\tilde{\sigma}_j$, $0 \leq \tilde{\sigma}_j \leq 1$, and γ_i is the random effect measuring the variability (or reliability) of the rating firm i . We assume that the γ_i arises from a common normal distribution with zero mean and common variance.

APPENDIX B: SIMULATION RESULTS

This appendix reports on a series of four simulations designed to demonstrate the operation of the proposed optimization and Bayesian methods.

Table 2. Data and results for Simulation 1

Rater	Rating for			
	Firm 1	Firm 2	Firm 3	Firm 4
Firm 1	NA	0.50	0.40	0.10
Firm 2	0.70	NA	0.30	NA
Firm 3	0.85	NA	NA	0.15
Firm 4	NA	0.50	0.50	NA
Truth	0.50	0.25	0.20	0.05
Opt pair	0.48	0.24	0.21	0.07
Opt individual	0.48	0.24	0.21	0.07
Bayes linear	0.49	0.21	0.21	0.09
Bayes comp	0.41	0.28	0.24	0.06

Simulation 1. Our first simulation example uses the data introduced in Part 4.2 and reproduced in Table 2. The authors constructed the dataset by taking the ground truth (obviously not observed by the model) and then playing the role of each of the rating firms and doing rough estimates of the other firms. So, for example, Firm 2 only rates Firms 1 and 3, whose contributions in truth should be in a 5:2 ratio, but which we made roughly 70/30. As seen in Table 2, the models do a reasonable job estimating the “true” contribution values.⁴¹ The compositional model is notably less accurate, and we suspect this is because of the limited data available (and relatively large number of parameters) in a four-firm problem. The posterior distributions of the estimates for the compositional model exhibit a lot of variances, probably due to overfitting. For example, as seen in Table 3, the credibility intervals for the compositional model’s estimates are much wider than those for the linear model. These issues (and the relative success of the other methods) suggest that future work might consider some kind of regularization for the compositional model when dealing with small numbers of firms.

Simulation 2. The second simulation uses the data in Table 4. The dataset was constructed similarly to Simulation 1, except that it involves a larger group of firms and therefore more ratings. We also informally increased the amount of measurement error and introduced more missing data values. With six firms, all of the models perform well at recovering the “truth.” Once again, the credibility intervals for the Bayesian compositional model’s estimates are much

41 The Bayesian models were estimated using MCMC methods using the Stan statistical modeling platform. Visual checks of the trace plots suggest that the linear model has difficulty mixing, likely because of the (unmodeled) correlated errors. The compositional model exhibits no such issues.

Table 3. Credibility intervals for Bayesian estimates in Simulation 1

	Credibility interval for			
	Firm 1	Firm 2	Firm 3	Firm 4
Linear (50 percent CI)	(0.49, 0.49)	(0.21, 0.21)	(0.21, 0.21)	(0.09, 0.09)
Linear (95 percent CI)	(0.41, 0.50)	(0.20, 0.26)	(0.20, 0.26)	(0.07, 0.09)
Comp (50 percent CI)	(0.22, 0.52)	(0.15, 0.42)	(0.14, 0.31)	(0.04, 0.08)
Comp (95 percent CI)	(0.01, 0.89)	(0.01, 0.85)	(0.01, 0.66)	(0.01, 0.17)

Table 4. Data and results for Simulation 2

Rater	Rating for					
	Firm 1	Firm 2	Firm 3	Firm 4	Firm 5	Firm 6
Firm 1	NA	0.60	0.30	NA	0.10	NA
Firm 2	0.50	NA	0.20	0.30	NA	NA
Firm 3	NA	NA	NA	0.50	0.25	0.25
Firm 4	NA	0.45	0.30	NA	0.15	0.10
Firm 5	0.50	0.50	NA	NA	NA	NA
Firm 6	0.60	NA	NA	0.35	0.05	NA
Truth	0.30	0.30	0.15	0.15	0.05	0.05
Opt pair	0.29	0.27	0.14	0.17	0.06	0.07
Opt individual	0.29	0.27	0.14	0.17	0.06	0.07
Bayes linear	0.28	0.28	0.11	0.17	0.08	0.08
Bayes comp	0.32	0.28	0.14	0.15	0.05	0.05

wider than those for the Bayesian linear model, as seen in Table 5,⁴² and visual checks of the trace plots yielded results similar to Simulation 1.

Simulation 3. The problem with manual encoding is that measurement errors are introduced in a haphazard and uncontrolled way. Thus, the “truth” is no longer really the truth, since our manual coding may inadvertently bias things in some direction. For the third simulation, we adopted a more systematic approach: We began with the “true” distribution among ten firms. Then for each rater, we added independent normal error to each component, where the standard deviation of the error was set at 20 percent of the true value (to model the fact that people are perhaps less precise when observing larger quantities). We then randomly dropped some of these components, although to make things more realistic, more “involved” firms (those entitled to a greater share) were more likely to observe a greater number of their peers. Finally, we

42 Hereafter, we will forego reporting credibility intervals in the interests of brevity.

Table 5. Credibility intervals for Bayesian estimates in Simulation 2

	Credibility interval for					
	Firm 1	Firm 2	Firm 3	Firm 4	Firm 5	Firm 6
Linear (50 percent CI)	(0.28, 0.30)	(0.28, 0.30)	(0.11, 0.15)	(0.15, 0.17)	(0.05, 0.08)	(0.06, 0.08)
Linear (95 percent CI)	(0.24, 0.32)	(0.25, 0.35)	(0.11, 0.18)	(0.10, 0.19)	(0.02, 0.08)	(0.05, 0.14)
Comp (50 percent CI)	(0.22, 0.41)	(0.18, 0.32)	(0.10, 0.18)	(0.10, 0.19)	(0.04, 0.07)	(0.04, 0.06)
Comp (95 percent CI)	(0.04, 0.71)	(0.05, 0.58)	(0.10, 0.40)	(0.10, 0.40)	(0.01, 0.12)	(0.01, 0.10)

renormalized the observations so that the relative shares all added to one. The resulting set of observations was labeled the “Baseline” set for the simulation.

To simulate collusive behavior, we then assumed that Firms 4 and 5 agreed to inflate each other’s scores. Table 6 shows an example of the baseline (no collusion) set of observations for Firms 4 and 5, and then the affected (collusion) set for the same two firms. As we emphasize in bold, as a result of the collusion, Firm 4’s rating for Firm 5 is about twice as large, whereas Firm 5, who would have not rated Firm 4 at all, instead rates Firm 4 greater than Firm 1, the greatest contributor in the set. Those inflated scores then have concomitant downstream effects on the other firms rated.

Table 7 displays the true contribution for each of the ten firms in Simulation 3, and the results of the various models on the Baseline and Collusion datasets. Three results are especially worthy of note. First, the outputs of the four models are again quite similar. Second, in all practicality, all of the methods are largely collusion resistant. Both Firms 4 and 5 deviated significantly from their honest reports, actions extreme enough to risk detection and possible sanction. Yet, as seen in bold in Table 7, at best, they are able only to affect outcomes by a few percentage points, often within the general noise we see in the table. The reason for this inherent collusion resistance likely comes from our method’s use of a “web” of peer assessments. With multiple firms rating multiple peers, no single rating has a significant effect on the outcome, no matter how extreme it may be. Third, among the various methods, the Bayesian collusion resistant model seems best able (at least in this example) to negate the distortion created by Firm 5’s collusive efforts.

One should bear in mind that all of these results are subject to random variation. We construct the baseline dataset through random processes, and the process of constructing the collusion set (which attempts to distort the baseline set) also has random aspects. Further, the estimation procedure for the Bayesian model, which uses MCMC methods, also has random aspects. To get a better sense of average model performance for the Bayesian models, we ran a procedure similar to Simulation 3 (dataset generation as well as model

Table 6. Excerpted baseline and collusion data for Simulation 3

	Rating for firm									
	1	2	3	4	5	6	7	8	9	10
Baseline 4	0.42	NA	0.28	NA	0.19	NA	0.11	NA	NA	NA
Baseline 5	0.36	0.26	0.21	NA	NA	0.17	NA	NA	0.079	0.083
Collusion 4	0.32	NA	0.21	NA	0.39	NA	0.08	NA	NA	NA
Collusion 5	0.21	0.16	0.13	0.39	NA	0.10	NA	NA	NA	NA

Table 7. Results for Simulation 3

	Allocation									
	1	2	3	4	5	6	7	8	9	10
Truth	0.25	0.15	0.15	0.10	0.10	0.10	0.05	0.05	0.025	0.025
Baseline set										
Opt pair	0.24	0.16	0.14	0.11	0.07	0.12	0.05	0.06	0.032	0.026
Opt individual	0.24	0.16	0.13	0.12	0.08	0.12	0.05	0.06	0.029	0.024
Bayes no resist	0.25	0.15	0.13	0.11	0.09	0.11	0.05	0.06	0.028	0.028
Bayes resist	0.23	0.16	0.13	0.11	0.09	0.10	0.05	0.06	0.032	0.029
Collusion set										
Opt pair	0.20	0.14	0.11	0.14	0.09	0.11	0.05	0.08	0.058	0.043
Opt individual	0.20	0.15	0.11	0.16	0.09	0.10	0.04	0.07	0.045	0.036
Bayes no resist	0.24	0.15	0.13	0.13	0.08	0.10	0.05	0.06	0.028	0.028
Bayes resist	0.25	0.14	0.13	0.11	0.07	0.10	0.05	0.06	0.027	0.027

estimation) twenty times, measuring the performance of the models by calculating the sum of square errors between the estimated and true parameters, namely:

$$\text{Performance Metric} = \sum_j (\hat{\alpha}_j - \alpha_j)^2$$

The results of this exercise are seen in [Table 8](#). It shows that the noncollusion-resistant model predictably performs less well on the collusion dataset than on the baseline dataset. This result is of course expected, since the colluders are distorting the observations. The collusion-resistant model, however, seems able to counteract some of the collusion, getting on average closer to the “truth” than the nonresistant model.

Simulation 4. To test how our methods handled self-reports, we extended Simulation 3 (specifically the simulation in [Table 7](#)) to include self-reports. We estimated the allocations for four sets of data, all based on the original baseline set from Simulation 3: (i) the baseline set as before; (ii) the baseline set plus

Table 8. Average mean squared error from twenty runs of Simulation 3

	Dataset/Bayesian model		
	Baseline/no resist	Collusion/no resist	Collusion/resist
Mean performance	1.49E-3	2.65E-3	1.98E-3

Table 9. Results for self-report simulation

	Rating for firm									
	1	2	3	4	5	6	7	8	9	10
Truth	0.25	0.15	0.15	0.10	0.10	0.10	0.05	0.05	0.025	0.025
Opt pair										
No self-report	0.24	0.16	0.14	0.11	0.07	0.12	0.05	0.06	0.032	0.026
Unbiased self-report	0.23	0.17	0.14	0.11	0.08	0.11	0.05	0.06	0.024	0.025
Firm 3 inflated self	0.23	0.16	0.16	0.11	0.08	0.11	0.05	0.06	0.026	0.027
Firm 9 inflated self	0.22	0.16	0.13	0.11	0.09	0.11	0.05	0.06	0.060	0.024
Opt individual										
No self-report	0.24	0.16	0.13	0.12	0.08	0.12	0.05	0.06	0.029	0.024
Unbiased self-report	0.23	0.16	0.14	0.12	0.08	0.11	0.05	0.06	0.025	0.025
Firm 3 inflated self	0.22	0.16	0.16	0.11	0.08	0.11	0.05	0.06	0.028	0.029
Firm 9 inflated self	0.22	0.16	0.13	0.11	0.09	0.11	0.05	0.06	0.053	0.023
Bayes no resist										
No self-report	0.25	0.15	0.13	0.11	0.09	0.11	0.05	0.06	0.028	0.028
Unbiased self-report	0.24	0.16	0.13	0.12	0.08	0.12	0.05	0.06	0.027	0.027
Firm 3 inflated self	0.23	0.15	0.14	0.10	0.07	0.12	0.05	0.06	0.027	0.027
Firm 9 inflated self	0.22	0.16	0.13	0.11	0.07	0.12	0.05	0.06	0.030	0.028
Bayes resist										
No self-report	0.23	0.16	0.13	0.11	0.09	0.10	0.05	0.06	0.032	0.029
Unbiased self-report	0.22	0.16	0.13	0.10	0.09	0.12	0.05	0.06	0.025	0.025
Firm 3 inflated self	0.26	0.17	0.13	0.10	0.06	0.12	0.05	0.06	0.025	0.025
Firm 9 inflated self	0.22	0.16	0.13	0.10	0.09	0.12	0.05	0.06	0.032	0.028

unbiased self-reports; (iii) the baseline set with self-reports where Firm 3 doubled its own rating; and (iv) the baseline set with self-reports where Firm 9 quadrupled its own rating. As seen in Table 9, the models appear unfazed by the addition of self-reports, whether unbiased or distortive.

REFERENCES

- Nfl settlement program summary report. Technical report, 2021. URL <https://www.nflconcussionsettlement.com/>.
- Aitchison John. *The statistical analysis of compositional data*. Monographs on statistics and applied probability (Series). Chapman and Hall, London, 1986. ISBN 0412280604.

- Alchian Armen A. and Harold Demsetz. Production, information costs, and economic organization. *The American Economic Review*, 62(5):777–795, 1972. ISSN 00028282. URL <http://www.jstor.org/stable/1815199>.
- Avriel Mordechai. *Nonlinear programming: analysis and methods*. Prentice-Hall, Englewood Cliffs, N.J, 1976. ISBN 0136236030.
- Baiman Stanley and Joel S. Demski. Economically optimal performance evaluation and control systems. *Journal of Accounting Research*, 18:184–220, 1980. ISSN 00218456, 1475679X. URL <http://www.jstor.org/stable/2490338>.
- Baldenius Tim, Jonathan Glover, and Hao Xue. Relational contracts with and between agents. *Journal of Accounting and Economics*, 61(2-3):369–390, 2016. ISSN 0165-4101.
- Becker Charles L., Shanin Specter, and Thomas R. Kline. How not to manage a common benefit fund: Allocating attorneys’ fees in Vioxx litigation. *Drexel Law Review*, 9:1–41, 2016.
- Belson Ken. Judge approves deal in N.F.L. concussion suit. *New York Times*, page B11, 2015. URL <https://www.nytimes.com/2015/04/23/sports/football/nfl-concussion-settlement-is-given-final-approval.html>.
- Steven Brams J. and Alan D. Taylor. *Fair division: from cake-cutting to dispute resolution*. Cambridge University Press, Cambridge, 1996. ISBN 0521553903.
- Bronstad Amanda. Critical mass: Class needs replacement as embattled Michael Avenatti exits. plus, fee battle roils mesh litigation, 2019a. URL <https://www.law.com/2019/04/03/critical-mass-class-needs-replacement-as-embattled-michael-avenatti-exits-plus-fee-battle-roils-mesh-litigation/>.
- Amanda Bronstad. ‘Structural and procedural flaws’ foul up fees in Syngenta settlement, 2019b. URL <https://www.law.com/nationallawjournal/2019/08/28/structural-and-procedural-flaws-foul-up-fees-in-syngenta-settlement/>.
- Burch Elizabeth Chamblee. *Mass tort deals: backroom bargaining in multidistrict litigation*. Cambridge University Press, Cambridge, United Kingdom, 2019.
- Che Yeon-Koo and Seung-Weon Yoo. Optimal incentives for teams. *The American Economic Review*, 91(3):525–541, 2001. ISSN 00028282. URL <http://www.jstor.org/stable/2677877>.
- Clermont Kevin M. and John D. Currivan. Improving on the contingent fee. *Cornell Law Review*, 63:529–639, 1977.
- Coffee John C. Understanding the plaintiff’s attorney: The implications of economic theory for private enforcement of law through class and derivative actions. *Columbia Law Review*, 86:669–727, 1986.

- de Clippel Geoffroy, Herve Moulin, and Nicolaus Tideman. Impartial division of a dollar. *Journal of Economic Theory*, 139(1):176–191, 2008. ISSN 0022-0531.
- Deb Joyee, Jin Li, and Arijit Mukherjee. Relational contracts with subjective peer evaluations. *The RAND Journal of Economics*, 47 (1):3–28, 2016. doi: <https://doi.org/10.1111/1756-2171.12116>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12116>.
- Eisenberg Theodore and Geoffrey P. Miller. Attorney fees and expenses in class action settlements: 1993–2008. *Journal of Empirical Legal Studies*, 7: 248–281, 2010.
- Epstein Richard A. Class actions: The need for a hard second look. *Civil Justice Report*, 2002.
- Fitzpatrick Brian T. The end of objector blackmail. *Vanderbilt Law Review*, 62:1623–1666, 2009.
- . An empirical study of class action settlements and their fee awards. *Journal of Empirical Legal Studies*, 7:811–846, 2010a.
- . Do class action lawyers make too little. *University of Pennsylvania Law Review*, 158:2043–2083, 2010b.
- . *The conservative case for class actions*. The University of Chicago Press, Chicago, 2019. ISBN 9780226659336.
- Gelman Andrew. Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian Analysis*, 1(3):515–534, 09 2006. doi: 10.1214/06-BA117A. URL <https://doi.org/10.1214/06-BA117A>.
- Gershkov Alex and Eyal Winter. Formal versus informal monitoring in teams. *American Economic Journal: Microeconomics*, 7(2):27–44, April 2015. doi: 10.1257/mic.20130277. URL <https://www.aeaweb.org/articles?id=10.1257/mic.20130277>.
- Gershkov Alex, Jianpei Li, and Paul Schweinzer. How to share it out: The value of information in teams. *Journal of Economic Theory*, 162:261–304, 2016. ISSN 0022-0531.
- Green Jerry R. and Nancy L. Stokey. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–364, 1983. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/1837093>.
- Holmstrom Bengt. Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340, 1982. ISSN 0361915X. URL <http://www.jstor.org/stable/3003457>.
- Holmström Bengt and Paul Milgrom. Regulating trade among agents. *Journal of Institutional and Theoretical Economics*, 146(1):85–105, 1990. ISSN 0932-4569.

- Itoh Hideshi. Cooperation in hierarchical organizations: An incentive perspective. *Journal of Law, Economics, and Organization*, 8(2):321–345, 1992. ISSN 87566222, 14657341. URL <http://www.jstor.org/stable/765083>.
- . Coalitions, incentives, and risk sharing. *Journal of Economic Theory*, 60(2):410–427, 1993. ISSN 0022-0531.
- Kandel Eugene and Edward P. Lazear. Peer pressure and partnerships. *Journal of Political Economy*, 100(4):801–817, 1992. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/2138688>.
- Klonoff Robert H. *Federal multidistrict litigation in a nutshell*. West Academic Publishing, St. Paul, MN, 2000.
- Landeo Claudia M and Kathryn E Spier. Incentive contracts for teams: Experimental evidence. *Journal of Economic Behavior and Organization*, 119: 496–511, 2015. ISSN 0167-2681.
- Lazear Edward P. and Sherwin Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864, 1981. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/1830810>.
- Ma Ching-to, John Moore, and Stephen Turnbull. Stopping agents from “cheating”. *Journal of Economic Theory*, 46(2):355–372, 1988. ISSN 0022-0531.
- Macho-Stadler Inés and J.David Pérez-Castrillo. Moral hazard with several agents: The gains from cooperation. *International Journal of Industrial Organization*, 11(1):73–100, 1993. ISSN 0167-7187.
- Mookherjee Dilip. Optimal incentive schemes with many agents. *The Review of Economic Studies*, 51(3):433–446, 1984. ISSN 00346527, 1467937X. URL <http://www.jstor.org/stable/2297432>.
- Moulin Hervi. *Axioms of Cooperative Decision Making*. Econometric Society Monographs. Cambridge University Press, 1988. doi: 10.1017/CCOL0521360552.
- Mitchell Polinsky A. and Daniel L. Rubinfeld. Aligning the interests of lawyers and clients. *American Law and Economics Review*, 5(1):165–188, 2003. ISSN 14657252, 14657260. URL <http://www.jstor.org/stable/42706198>.
- Resnik Judith. Vital state interests: From representative actions for fair labor standards to pooled trusts, class actions, and MDLs in the federal courts. *University of Pennsylvania Law Review*, 165:1765–1806, 2017.
- Robertson Jack and William Webb. *Cake-Cutting Algorithms: Be Fair if you Can*. AK Peters/CRC Press, Natick, MA, 1998.
- Rubenstein William B. On what a common benefit is, is not, and should be. *Class Action Attorney Fee Digest*, 3:87–94, 2009.
- Rubenstein William B. and Herbert B. Newberg. *Newberg on class actions*. West/Thomson Reuters, Eagan, MN, 5 edition, 2011.

- Segal-Halevi Erel. Fairly dividing a cake after some parts were burnt in the oven. <https://arxiv.org/abs/1704.00726>, 2018.
- Shavell Steven. *Foundations of economic analysis of law*. Belknap Press of Harvard University Press, Cambridge, Mass, 2004. ISBN 0-674-04349-9.
- Siemens Enno, Sridhar Balasubramanian, and Aleda V. Roth. Incentives that induce task-related effort, helping, and knowledge sharing in workgroups. *Management Science*, 53(10):1533–1550, 2007. ISSN 00251909, 15265501. URL <http://www.jstor.org/stable/20122312>.
- Silver Charles and Geoffrey P. Miller. The quasi-class action method of managing multi-district litigations: Problems and a proposal. *Vanderbilt Law Review*, 63:105–177, 2010.
- Stan Development Team. *Stan Modeling Language Users Guide and Reference Manual*. 2020.
- Tideman T. Nicolaus and Florenz Plassmann. Paying the partners. *Public Choice*, 136(1/2):19–37, 2008. ISSN 00485829, 15737101. URL <http://www.jstor.org/stable/27698288>.
- Varian Hal R. Monitoring agents with other agents. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, 146(1):153–174, 1990. ISSN 09324569. URL <http://www.jstor.org/stable/40751313>.