Valuation Follies

Andrew C. Baker[*] Jonah Gelbach[†] Eric Talley[‡]

Abstract

Financial valuation has become central to the resolution of legal disputes across substantive areas of commercial litigation. Ostensibly grounded in academic theory, valuation for litigation purposes is frequently untethered from contemporary practice in finance. In this paper we show how the conventional approach to valuation used in litigation suffers from a number of conceptual flaws, which results in substantial discretion on behalf of economic experts providing valuation estimates. This discretion frequently frustrates generalist judges, who are forced to adjudicate complicated and subjective disputes outside of their area of comfort. In this paper we propose alternative, data-driven approaches to valuation that largely avoid the pitfalls of the conventional approach. We show through a simulation analysis that our proposed alternatives provide more accurate and lower-variance estimates.

[*]University of California, Berkeley
[†]University of California, Berkeley
[‡]Columbia University

1 Introduction

For the past four decades, financial valuation has played an increasingly pivotal role in the litigation of high-stakes commercial disputes. While the judicial embrace of such methodologies was initially limited to isolated topics in corporate and securities law, the practice quickly expanded. Litigation has come to be influenced, and often dominated, by valuation disputes that hinge on financial economics—from bankruptcy to tax disputes, family law, fiduciary duties, and garden-variety questions in tort, property and contract law. By some accounts, the incursion of modern finance into commercial law has been nothing short of a pioneering "revolution"–a long-overdue hostile takeover of an "ossified, stagnant field". Every top US law school now offers at least one course dedicated to teaching these techniques to law students.

However, the wholesale adoption of financial valuation in commercial litigation smuggled in a hidden adversary of its own making, and as modern finance infused substantive law, it unleashed at least four undesirable collateral consequences. The first stems from the fact that financial economics tends to be a mathematical enterprise. As quantitative methodologies found their way into litigation settings, they quickly confronted courts with a demanding progression of technical challenges. Most judges are not formally trained asset-pricing specialists, even in business-intensive courts like Delaware's Court of Chancery; yet they are now typically required to admit, exclude, and sometimes weigh technical financial evidence, or even instruct lay juries possessing even less expertise. In most cases, judges are left to pick up key tenets of valuation practice on the fly, frequently (and understandably) relying on the motivated pedagogy of litigants or their hired experts.

Second, despite its seemingly precise technical façade, there is little doubt that financial valuation as practiced is as much art as science: the field is awash with free parameters that both require judgment and afford considerable discretion to the expert analyst. Not only can an expert select from a sizable menu of general approaches to render estimates of fair market value, but within each lurks a number of dubious and subjective assumptions around implementation. In addition, the professional literature that purportedly guides and substantiates financial experts' choices has itself grown heterogeneous, offering a variety of self-styled "authoritative manuals" that contain distinct—and surprisingly inconsistent—formulations for best practices, further amplifying the need for judgment (and the discretion that such judgment entails).

Third, those deploying financial methodologies in the courtroom are typically the litigants' own compensated experts. While most are respected members of academic and/or professional finance circles, they tend to be repeat-player consultants who perform their work largely shielded from public scrutiny, crafting reports that can cherry-pick from assorted best-practice formulations in ways that—(un)remarkably—favor their clients' economic interests. Less self-serving (and more moderate) calculations may be left on the cutting room floor. Expert reports, moreover, are often filed under seal, remaining secluded from public view until long after trial, perhaps indefinitely. With no real threat of professional scrutiny, little stops dueling experts from embracing techniques that they would eschew or even deride in academic settings, resulting in valuations that not only may be inconsistent with academic practices, but also frequently diverge by orders of magnitude.

Finally, in groping to cope with an already-flawed valuation ecosystem, courts have unwittingly distorted it further. Non-expert judges who adjudicate the claims of dueling financial experts must provide reasons for their judgments. That reasoning frequently adopts one or the other expert's assumptions for each specific issue, though it might also split the difference between experts. Either way, the judge's reasoning—especially once memorialized in a written decision—can implicitly entrench and amplify the discretionary choices made by experts in generating best practices. Over time, this tapestry becomes less financial economics and more a "rulified" distillation of judicial folk-wisdoms about finance—one whose core tenets wander afield from contemporary social science. This assumes a life of its own, with each successive opinion contributing to a self-perpetuating echo chamber of accepted conventions. The end result is something resembling a distinct (and ironically nonacademic) legal doctrine, which—while sprouting from the finance scholarship of the 1970s and 1980s—has subsequently followed a markedly different path, impelled substantially by interventions from a coterie of motivated litigators, experts, litigation consultants, and best-practice entrepreneurs.

In recent years, various commentators have offered suggestions for how to best remedy the dysfunctional system described above. For example, at the urging of several academic commentators, some courts have increasingly embraced the practice of routinely unsealing expert reports for public scrutiny (at least after trial), under the theory that the prospect of professional embarrassment can provide needed discipline. Others have suggested structural reforms, such as having courts retain an independent expert to advise on valuation issues, or experimenting with expert "hot-tubbing", or committing to final-offer (a.k.a. "baseball") arbitration mechanisms to incentivize experts towards moderation. Still others have instead attempted to skirt altogether the messy enterprise of valuation sausage making, by, for example, simply adverting to the negotiated merger price itself or pre-existing securities market prices.

Although we believe that many of these institutional tweaks warrant consideration, this project follows a different path: We advance the thesis that now is an opportune moment for courts to reconcile courtroom practice with recent academic advances, which are increasingly being fueled by machine learning. Our argument starts with a simple observation: Notwithstanding any technical trappings, every valuation methodology in use is ultimately an exercise in prediction. Whether used to estimate the value of a firm's equity, debt, enterprise value, or other target, all valuation methodologies and "best practices" are worth deploying only because they are thought to render a credible prediction of that target given the available data. In fact, each of the three leading market valuation methodologies in use today—comparable transactions analysis, comparable companies analysis, and discounted cash flow analysis—aspires to deliver a prediction of firm value; the three methodologies differ not in this ultimate goal, but rather in the choice and use of data to accomplish the task.

Given this reality, it follows that the science of prediction becomes a critical reference point for our collective attention. In this paper, we argue that simple, off-the-shelf machine learning techniques would provide interpretable results that outperform current practice, and also that they are perfectly consistent with the doctrines and evidentiary rules that govern litigation.

Our project therefore makes three contributions. First, using simulation evidence based on actual firm valuations, we show that current practices allow considerable expert discretion—what we refer to as "expert degrees of freedom." Second, we argue that emerging tools from machine learning (ML) may be able to augment, improve, or even supplant prevailing courtroom valuation practices. Third, we argue not only that judges can admit expert evidence based on such approaches, but also that they can and should begin to demand them from litigants and experts.

Our simulation evidence regarding expert degrees of freedom appears in section III has one overarching finding: conventional valuation practices, which we argue are essentially nearest-neighbor algorithms, bring with them substantial expert degrees of freedom, even when only minor variations in the parameterization of the nearest-neighbor routine are allowed.

Regarding our second contribution, we demonstrate that simple statistical learning methods can be used to substantially improve the predictive performance of conventional valuation methodologies. We start by observing that each of the three primary approaches to valuation exhibit clear similarities to primitive ML algorithms. Accordingly, each might, in principle, be improved by disciplining current practice using only rudimentary improvements. In Section III, our results suggest that a simple data-driven approach to valuation clearly outperforms a random choice from among the 24 parameterizations we consider.

We acknowledge that these ideas are not without challenges. Machine learning techniques do best, all else equal, when large quantities of comparable data are available, and that is not always true in valuation disputes. Still, the large amount of publicly available financial data will make it feasible in many instances. And insofar as current practices are essentially rudimentary nearest-neighbor algorithms, critiques of ML feasibility are also critiques of current practice.

Our third contribution is to argue that ML-based expert evidence is perfectly consistent with current evidence requirements in effect in state and federal courts. Federal Rule of Evidence 702, for example, conditions expert testimony on establishing the expert's qualifications in "knowledge, skill, experience, training, or education." The Rule requires an expert's testimony to be based on "sufficient facts or data," and also that the testimony is "the product of reliable principles and methods" that are "reliably applied ... to the facts of the case." Because Rule 702's 2000 amendment codified *Daubert*, the Rule embraces *Daubert*'s flexible standard concerning the testability of the expert's opinion—its academic pedigree, error rate, and acceptance within a relevant scientific community. The majority of states, including Delaware, follow *Daubert*'s principles, and Delaware's Supreme Court has applied them in numerous cases. That said, certain states continue to follow the pre-*Daubert Frye* standard that "the thing from which the deduction is made must be sufficiently established to have gained general acceptance in the particular field in which it belongs."

Data science and machine learning, we argue, are not merely promising novelties that may, someday, satisfy these admissibility standards. Rather, the practice is already established, having in the past two decades spawned numerous peer reviewed journals, academic departments, and scholarly conferences worldwide. In fact, we maintain that the data-driven predictive tasks that ML-based models are designed to solve transparently deliver precisely the kind of reliability diagnostics that would satisfy Rule 702 and *Daubert*. And, because such models are designed to optimize direct statistical measures of prediction accuracy, they often will permit side-by-side comparisons of experts' competing models. By their construction, moreover, well-constructed ML-based methods will typically do no worse than the standard approaches, and, given enough data, will often be expected to fare better. Consequently, should courts take up our invitation to embrace ML-based valuation approaches, traditional cookie-cutter valuation methodologies may well begin to face increasingly difficult admissibility challenges of their own.

An important caveat to our analysis deserves mention before proceeding. As noted above, the dysfunction afflicting courtroom valuation at present stems in large part from the significant discretion that current valuation methodologies afford experts. Although our proposed ML-based methodology neutralizes many dimensions of discretion, we should be mindful that it—like any new methodology—may spawn a new set of conventions, norms of judgment, and dimensions of discretion for future experts to apply. It is therefore fair to question whether, if embraced, our proposed approach would eventually fall prey to a similar set of forces that have undermined the status quo. While we take this challenge seriously, we believe that the promise of data-driven predictions, even as a baseline comparator, deserves serious consideration.

Our analysis proceeds as follows. In Section II we provide an overview of standard valuation approaches that are prevalent in the literature, focusing on comparable transactions, comparable companies, and discounted cash flow analyses. Section III argues that the first two can be interpreted as a nearest-neighbor ML algorithm, with as-practiced discounted cash flow analysis sharing that feature at least partially. More directed ML approaches could both loosen those conventions and admit a richer variety of models to predict firm value. In Section IV, we spotlight comparable companies analysis and demonstrate the prevalence of expert degrees of freedom through the use of a simulation exercise. We propose and test a series of competing approaches that use data-driven methods to optimally estimate firm valuation from observable data. In Section V we provide an example of how our alternative ML-based models would work in an actual case, and we end with a discussion of some limitations and extensions of our approach in Section VI.

2 Conventional Practice

As it is currently practiced in business law courtrooms and boardrooms, modern valuation practice is dominated by three alternative methodologies: Comparable companies (CC), comparable transactions (CT), and discounted cash flow (DCF) approaches. In many cases, a valuation expert will attempt to value a financial asset of interest using two, or even all three of these approaches. Particularly in cases of company valuations associated with merger agreements or bankruptcies, experts may use other valuation methodologies as complements. Such additional approaches include an analysis of historical premiums paid, analyst forecasts, and leveraged buyout/recapitalization analysis. When such alternatives are employed, however, they are typically offered only as "reference" valuations meant to complement CC, CT and DCF approaches.

2.1 Comparable Transactions

The CT approach may be the most intuitively accessible, as it bears resemblance to the approach that real estate appraisers take when using "comps" to estimate the value of one's home. The basic idea is to find examples of analogous assets that have recently been sold in arm's length-transactions, and use those sales prices to deliver an estimate of what the sale of the company in question would deliver. With home appraisals, this process usually begins by selecting neighborhoods in a similar geographic area with similar traits, such as walkability, school quality, and income, as well as having a similar number of bedrooms and square footage. The recently sold properties deemed similar to the property in question along these dimensions are an appraiser's comps. The appraiser then will usually normalize the measure of value, such as price per square foot, for each comp transaction, and will aggregate the comps using the mean or median of the price-per-square-foot values. This yields a summary measure of the price per square foot to be applied to the property whose valuation is in question.[1]

The CT approach for companies and other financial assets operates similarly. Much like the property appraiser, a valuation analyst using a CT methodology will first find recent sales of companies deemed comparable. If feasible, the analyst will limit attention to transactions in similar industries, geographic locations, or vintages.[2] Like real estate, companies vary in size, so finding comps of similar scale is desirable. In addition, they can have unique capital structure traits: for example, corporate debt often can transfer over as part of the sale, in which case the purchase price reflects only equity value. The valuation analyst will thus attempt to produce a measure of firm value that controls for both equity and debt factors. To control for capital structure, the analyst often needs to rescale the purchase price to reflect what is known as the "enterprise value" of each comparable company, adjusting the sales price of each comp to account for the value of any debt not capitalized into the sales price.[3]

Once comparable sales prices are converted to enterprise values, analysts then address the size factor, using an analog of the price per square foot measure used by real estate appraisers. Here, however, the standard normalized metric is typically an earnings multiple: That is, re-expressing the enterprise value not in raw dollar terms, but rather as a multiple of some specified measure of earnings. A standard

[1] Either a point estimate or a range of estimates might be provided.

[2] These and other traits are specified in a small set of valuation manuals that have become accepted over time in the profession.

[3] Other adjustments include netting off cash (and cash equivalents), as well as making sometimes-controversial changes to working capital.

metric that operates as a default is earnings before interest, taxation, depreciation and amortization (EBITDA), which is often a relatively stable proxy for cash flows in mature companies. For less mature companies, however, it is not uncommon to see multiples based on other factors, such as EBIT, sales revenues, or, less commonly, other measures of market interest such as "clicks" on the company's website. Non-EBITDA multiples are typically disfavored,[4] however, and tend only to be used when the company generates negative adjusted earnings, which renders any multiples-based approach nonsensical.[5]

Even after a multiple has been selected, there are many ways to quantify the denominator of the multiple. For example, one might base a multiple on the last fiscal year's numbers, or the last twelve months, or projections for the next twelve months or fiscal year. In each case, the multiple of the comparable company may change, and it is not uncommon for analysts to assess CT multiples using several measures, thereby cobbling together a range of valuations based on the aggregated outcomes of such approaches, as well as a variety of permutations of the mean, median or inter-quartile range for each measure. In formulating the comps, the multiple formulations, and the ranges, the analyst typically retains significant discretion—an issue to which we return in our analysis below.

Two potential constraints on the CT approach often limit its ability to deliver valuation projections. The first is a lack of data. Because bona fide arms-length sales of companies within a given industry are generally rare, the set of "comparable transactions" might itself be relatively sparse, which may force the analyst to make a projection from a very small group (perhaps even as small as one). One potential solution to this problem is to lengthen the time horizon or broaden the criteria by which comparatives are drawn (e.g., by expanding the industries considered). The second potential limitation on the CT approach is that it is predicated on sales of comparable firms through negotiated transactions. Such sales often come with a premium baked into the sales price, reflecting the value of control. This baked-in control premium may sometimes be inappropriate if (for example) one is interested in gauging only the cash flow valuation of the company. In such settings, CT requires an attempt to shave off control premia from precedent sales.

[4]Our simulation analysis focused on an EBIT-based measure for data availability reasons.

[5]We note that this is not a rare occurrence; in our data sample for the simulation, nearly 30% of firms have negative EBITDA values. Thus, we focus on EBIT in scaling earnings to increase sample size in our analyses.

2.2 Comparable Companies

The CC approach is a close cousin to the CT approach, differing only in the source of the data used to assess comparable firms. While CT uses sales price data from acquisitions of comparable firms, CC uses the full spectrum of data from large and thick securities markets. In thick markets, stock prices are thought to be a good proxy for the economic value of a fractional share of the company, at least on average and as viewed by the marginal investor. Thus, rather than using the sales price to predict value, the total market capitalization of comparable firms can be based on public trading data. Beyond that, however much of the CC valuation process is identical to that in CT, including the conversion from an enterprise value multiple to enterprise value, the specification of the multiple itself, the use of judgment about how to measure such multiples, such as last twelve months, next twelve months, etc., and a summary measure (mean or median) for aggregating comp multiples. In other words, beyond the different source of valuation metrics for the comparable firms, nearly every other part of a CC analysis tracks the CT analysis almost directly.

CC methodologies have one obvious advantage over CT approaches: data. It can be difficult to find precedent M&A transactions to use in developing CT comps, but CC is facilitated because thousands of public companies trade continuously and have observable prices each day. Consequently, CC allows one to build sizable groups of comparable firms. On the other hand, CC also tethers the value of companies to the trading value of their stocks, which in turn tends to reflect the value the market ascribes to a publicly held valuation target. Thus, the CC approach might neglect the value of the control premium. Another potential limitation of the CC approach is that it depends critically on the on-average value efficiency of trading markets. This may often be appropriate, but CC fits less well when the target firm is traded in a thin, volatile, poorly-developed market where market valuations and fundamental valuations can diverge.

2.3 Discounted Cash Flow

The third major form of valuation is discounted cash flow (DCF) analysis. Compared to CC and CT approaches, the DCF has significantly more moving parts, and is generally viewed as more technically demanding. Rather than looking for comparable firms (at least directly), the DCF approach conceives of the value of a financial asset as the equivalent to the present discounted value of the free cash flows the asset is projected to produce. Borrowing from the well-known formula in finance for present values, the DCF approach can be captured as follows:

$$FMV = PV(CashFlows) = \frac{FCF_1}{(1+WACC)} + \frac{FCF_2}{(1+WACC)^2} + \ldots + \frac{FCF_T}{(1+WACC)^T} + \frac{S_T}{(1+WACC)^{T'}}$$

where projected future time periods are divided into discrete units through some "terminal" projection period ($t = 1, 2, 3, \ldots T$); FCF_t denotes expected future "free cash flows" in each of the periods projected (typically 3-10 years into the future); S_t denotes a terminal (or "salvage") valuation of the asset as of the terminal projection year; and WACC represents a risk-adjusted discount rate known as the "Weighted Average Cost of Capital." If all of these ingredients are known (or can be reliably estimated under different scenarios), then a cash-flow prediction (or predicted range) of asset valuation is possible.

Like an onion, each of the ingredients of the DCF approach has its own layers of complexity (and resulting expert discretion). Free cash flows are sometimes generated from analyst forecasts, or through management forecasts done in the ordinary course, or through investment banker forecasts for the purposes of a (disputed) transaction, or by some other source. They are typically, though not always, unlevered, and thus do not carve off interest payable to capital creditors, so as to summarize the entire pool of earnings available to satisfy both debt holders and equity holders. And, in some cases cash flow projections of comparable firms (if available) can be used to form composite projections that more closely track the industry at large.

The WACC discount rate is typically a blend of expected return estimates for debt and equity (adjusted for leverage ratios and tax deductions), with equity return estimates the product of an underlying asset pricing model (such as the still-dominant market "beta" from the capital asset pricing model). In many cases, peer company betas are also blended in with company-specific data to create more of a composite measure (usually after an elaborate process adjusting peers' betas for differences in peer leverage ratios).

Finally, the terminal value measure (S_T) represents something of a capitulation to our inability to make projections indefinitely into the future. Because company projections are typically no longer than 10 years (and are more frequently 5-7 years), an analyst using DCF must make an assumption of what the asset will be worth in its terminal period (when no more forward looking projections are available). The assessment of terminal value is perhaps the most susceptible to expert degrees of freedom, since there is little to tether it to firm-level data. One approach for terminal values is to extrapolate the final period's free cash flow indefinitely into the future as a "growing perpetuity" at some posited growth rate, backing out a present valuation (as of period T) from a well-known formula for the present value of growing perpetuities.[6] Another frequently used approach, however, is simply to revert (once again) to peer company multiples using a recycled CC or CT approach applied as of the terminal period.

3 The Conventional Approach as K-Nearest Neighbors Matching

In Section 2 we described the conventional approach to valuing firms for litigation purposes. Experts retain substantial discretion in how they go about delivering a valuation—discretion that ranges from the choice of approach to how the measure the inputs of their valuation metric. Second, peer-group comparisons are pervasive—indeed such comparisons are the very backbone of CC and CT analysis, and even in DCF analysis peer group comparisons sneak in in myriad ways (free cash flow projections, terminal values, beta estimates, etc.).

When considering the scope of the expert's task, it is noteworthy how closely valuation practice resembles the k-nearest neighbor (k-NN) algorithm, both generally and in specific application. From a general perspective, conventional valuation measures reflect the same philosophy as machine learning generally, and the nearest-neighbor approach in particular, along three dimensions: (1) The goal is to deliver core predictions (rather than to test causal theories); (2) they use a significant amount of unstructured data to deliver those predictions; and (3) they habitually make use of claimed similarities of instances (via "comps") to buttress, support, and even drive their predictive enterprise.

The k-NN algorithm was first developed over a half century ago,[7] and is a non-parametric method for statistical learning often used to classify and/or make predictions about an outcome variable of interest by reference to the attributes of a given number, k, of the target's closest neighbors in some feature space, using a specific proximity measure. The method can be used for both classification problems, where the object is to predict which of a discrete set of categories the item of interest belongs, or regression problems, where the objective is to predict the value of a potentially continuously measured outcome variable's value. Valuation practices are

[6] For a posited perpetuity growth rate g, the formula is given by $S_T = \frac{FCF_{T+1}}{(WACC-g)} = \frac{FCF_T \times (1+g)}{(WACC-g)}$.

[7] See, e.g., Evelyn Fox & Joseph Hodges, (1951). Discriminatory Analysis. Nonparametric Discrimination: Consistency Properties Report). USAF School of Aviation Medicine, Randolph Field, Texas; Thomas M. Cover & Peter E. Hart (1967). "Nearest neighbor pattern classification" (PDF). IEEE Transactions on Information Theory. 13 (1): 21–27.

an instance of the latter, with the outcome variable being a measure of firm-value (either total market capitalization or enterprise value).

When using k-NN, the objective is to use observable data to predict an outcome label $\hat{y}$ for a new observation that has observed covariate measures X_0. In the simplest setting there is only one other variable x_0 used in the matching, and the k-nearest neighbors are the k units with the closest values of x_0 to the target unit. We then predict the outcome variable y for the target as the average of the k-nearest neighbors (with the average being either the mean or the median of the comparable units).

The k-NN algorithm can be generalized by adding large numbers of x-variables, as well as vector-valued outcomes. Its key ingredients are (a) an appropriate outcome variable y (often called a "label" in machine learning communities) for prediction; (b) a specified set of characteristics other than the outcome label by which to assess similarity/proximity (the xs); (c) a distance metric for measuring proximity (such as Euclidean distance); (d) a specified inclusion/weighting system for selecting comparators into the prediction set (such as picking the closest 3 neighbors); and (e) a metric for generating predictions from the included training variables (e.g. the mean, median, or inter-quartile range of the comparison-unit labels).

With respect to Comparable Companies and Comparable Transactions analysis, this discussion reveals that the two valuation methodologies are not simply similar to the k-NN algorithm—they are the kNN algorithm, if implemented in a somewhat casual form. Both approaches involve (a) taking earnings multiples of either trading values (CC) or acquisition values (CT) as the relevant outcome variable; (b) a specified set of data for identifying comparable companies/transactions; (c) a comparability assessment (although in practice this is rarely specified precisely); (d) rules of thumb for determining the number of comps to use (also involving expert discretion); and (e) a means for predicting the earnings multiple of the company in question (often the mean or median of the comparable firms, again at the expert's discretion). Indeed, while neither the CC method nor the CT method has to our knowledge previously been directly identified with k-NN algorithm, their resemblance is evident.

Although DCF valuation does not neatly map into the k-NN algorithm to the same degree, it, too, shares some features with the process. For example, it is common to use comparables to generate estimates of terminal value, earnings projections, and asset betas—all of which are a core ingredients of the DCF model. Thus, the spirit of the k-NN approach often enters DCF analysis at several junctures in material ways.

But if standard valuation methodologies are functionally k-NN learners, it is important to understand the approach's potential limitations. The k-NN approach can be attractive for a variety of reasons. First, k-NN is a form of instance-based learning, which means that, unlike other supervised learning approaches, it does not require a training stage for use. This makes k-NN simpler and faster to use than other algorithms that require training. Finally, when the data set grows large, k-NN regressions are known to be Bayes optimal.

On the other hand, k-NN classification has several limitations that can hamper its performance. First, as the amount of data grows, the cost of calculating distances increases (this is an example of the so-called "curse of dimensionality" that many non-parametric methods confront). k-NN use is also complicated by the presence of high-dimensional data (i.e. many xs), although this is not a problem with the number of variables typically used in valuation. Third, k-NN can be highly sensitive to scale and how distance is calculated, so good applications of k-NN must involve normalization before applying the algorithm—which adds a layer of expert discretion, because multiple scalers exist. Finally, k-NN can be sensitive to noisy data, missing values, and outliers.

4 Monte Carlo Simulation

In this section, we conduct a simple Monte Carlo simulation on actual firm valuation data to test how data-driven approaches to valuation compare with the conventional approach (in this case the CC method).

The majority of the data used in the simulation come from Compustat (financial reporting information) and the Center for Research in Security Prices (CRSP) (stock price data). We use quarterly financial data from 2000 to 2020 for all public reporting companies in the United States with a fiscal year reporting date of December 31,[8] which is subsequently merged with CRSP stock price data through the historical linking file provided by Wharton Research Data Services (WRDS). We keep only confirmed links,[9] and match to daily individual security prices and index returns from the CRSP daily stock and index files, as well as the daily factor returns from Ken French's website (i.e. the "Fama-French-Carhart Factors").

We next create a series of covariates for determining appropriate peers for valuation purposes based on the comparable considerations listed in leading Corporate Finance textbooks. A list of the firm-quarter variables, and their calculation method,

[8]We restrict our attention to 12/31 firms so that the different fiscal-year quarter ends align in calendar time. In unreported results we find that the inferences drawn from the analysis in this paper are unaffected by this choice.

[9]These are link codes equal to "LC", "LU", or "LS".

is provided in Appendix A. We require that there be non-missing entries for all of the covariates for a given observation to enter the simulation.

In addition, we categorize each firm-quarter observation as belonging to an industry, defined by the first two digits of the firm's Standard Industrial Classification (SIC) code. The primary SIC entry in Compustat (*sic*) is "header" information, which is the last identified observation for a given firm identifier and is static over time. We use the historical identifier (*sich*), which allows for time variation in industry designation. We impute missing entries of *sich* using a "down-up" strategy, which assumes that all missing values before the first recorded non-missing entry are equal to the first entry, and that subsequent missing values are equal to the most recent entry until a new entry is recorded. The header entry is used for missing entries after the last-in-time recorded industry designation. In this simulation we drop all observations for firms in the financial services industry (SIC code beginning with 6), because they do not have comparable values of sales and revenue.

With the constructed panel dataset of firm-quarter observations, we randomly sample 10,000 observations for benchmarking purposes. In order to be selected, we require that a given firm-quarter observation has i) at least nine peer firms with full non-missing data in the selected quarter, and ii) at least eight consecutive non-missing values for the ratio of market capitalization to earnings before interest and taxes (EBIT) ending on the selected quarter. In addition, for each sampled observation, we choose a random number of trading days, d^*, between zero and ninety, after the end of the quarter to value the firm. So, for example, if the randomly selected observation is for the third fiscal quarter of 2015, which ends on September 30, 2015, and we draw d^* = forty-five, we compare different valuation estimates for December 3, 2015, the forty-fifth trading date following quarter end.

For each draw in the Monte Carlo simulation, we have the observed firm valuation (i.e. market capitalization), a valuation ratio (the ratio of market capitalization to EBIT), an industry categorization, and a set of covariate values. We then compare how much the estimated valuation from various methods differ from the actual realized market valuation in the data. As a motivating example, we use one randomly-sampled observation for demonstrative purposes in the sections below:

Table 1: Motivating Example

Company	Quarter	Industry	Market Cap (M)	EBIT	Ratio	d^*
CVS HEALTH CORP	2010Q4	59	\$47,392	\$1,761	26.91	7

4.1 Conventional k-NN Approach

In Section 3 we described how the conventional approach to valuation described in leading corporate finance textbooks represents an applied form of k-nearest neighbor matching. In this section we use our simulation approach to test the performance, and susceptibility to expert discretion, of the conventional approach on realized financial outcomes.

The implementation of the conventional approach with our Monte Carlo sample proceeds in the following steps. For each randomly-chosen firm-quarter observation, we identify the target firm's quarterly valuation ratio r_{lq} for the target firm J in the target quarter q, which is the ratio of the firm's market capitalization at the end of the target quarter to the company's earnings before interest and taxes (EBIT) in that quarter. In addition, we calculate the market capitalization mc_{Jd^*} on the trading date d^* periods from the end of the quarter, and corresponding ratio of mc_{Jd^*} to the EBIT value for the target quarter q, r_{id^*}.[10]

We next identify all viable peer firms for the target company in that quarter. These are firms that are in the same two-digit SIC code industry and which have non-missing entries for the covariates.[11] For our motivating example, the potential peer firms for CVS Health Corp. are listed in Table 2.

Table 2: Potential Industry Peer Firms

MEDCO HEALTH SOLUTIONS INC	AMAZON.COM INC
OMNICARE INC	BIG 5 SPORTING GOODS CORP
CASH AMERICA INTL INC	QURATE RETAIL INC
FIRSTCASH HOLDINGS INC	HSN INC
TANDY LEATHER FACTORY INC	BLUE NILE INC
PCM INC	CABELAS INC
GLOBAL INDUSTRIAL CO	

This table reports the potential peer firms for our motivating example. CVS Health Corp. had a two-digit SIC code of 59 in Q4 2010, which is Miscellaneous Retail. The 13 peer firms in this table are those with the requisite data over the time period.

With the target firm and the potential peers, we next explore how to map the

[10]Conceptually, you can think of this as trying to value a firm sometime between fiscal quarter ends.

[11]For practical purposes we also require that the peer firms have full market trading data for the 250 trading days prior to quarter end, and the d^* days following the quarter.

input uncertainty into different valuations using k-NN. As mentioned earlier, there are multiple areas of discretion that can be used by an analyst when deriving a prediction from a k-NN algorithm. Here we focus on four: i) which variables to use for assessing the similarity of firms; ii) how many firms to match to; iii) how to measure the proximity of peer firms; and iv) whether to take the mean or median of the matched firm ratios when deriving a single summary parameter.

We first show how this works with a demonstrative case for our motivating example. We measure the similarity of the peer firms in Table 2 to CVS using the matching variables specified in the Rosenbaum & Pearl textbook. In Table 3 we rank the peers on similarity using the scaled euclidean distance between the target and each of the thirteen peer firms. After ranking the firms, we would then take, e.g., the arithmetic mean or median of the ratios for the top N matches, and use that ratio to impute the predicted market capitalization of CVS using target quarter EBIT of \$1.76 billion. If we were to take the median market cap to EBIT ratio values for the top 5 firms (40.21), it would imply a market capitalization for CVS on relative date d^* of approximately \$71 billion.[12]

Table 3: Motivating Example: Best and Worst Nearest Neighbor Matches

Company	EBIT	Market Cap (M)	Ratio	Rank
Best Matches				
MEDCO HEALTH SOLUTIONS INC	657.70	$27,162	41.30	1
AMAZON.COM INC	503.00	$83,137	165.28	2
BIG 5 SPORTING GOODS CORP	8.96	$320	35.73	3
OMNICARE INC	74.04	$2,977	40.21	4
HSN INC	77.44	$1,667	21.52	5
Worst Matches				
BLUE NILE INC	9.22	$827	89.71	9
TANDY LEATHER FACTORY INC	2.58	$48	18.71	10
FIRSTCASH HOLDINGS INC	27.47	$970	35.31	11
QURATE RETAIL INC	396.00	$9,097	22.97	12
PCM INC	6.80	$94	13.74	13

This table reports the five best and five worst matches for our motivating example using a k-nearest neighbors matching approach. We use the covariates from Rosenbaum and Pearl (RP), and a scaled euclidean distance metric.

[12] You back out the target market capitalization on date $d^* = 7$ by multiplying the matched peer ratio of 40.21 by CVS's 2010 Q4 EBIT value of \$1.76 billion, which equals approximately \$1.76 billion).

The set of matching variables, number of peers to match to, proximity distance metric, and summary measure were choices. In practice, an expert has considerable latitude to set these inputs differently in a manner that is perfectly consistent with the textbook descriptions of valuing firms. To explore the impact that different design choices can have on the valuation estimates, we estimate firm value using 24 different combinations of inputs:

- Two choices of matching variables (those described in either Rosenbaum & Pearl (RP) or Pratt & Niculita (PN)).
- Three choices for the number of matched firms (5, 7, or 9).
- Two different distance metrics (scaled euclidean (SE) or Mahalanobis (M)).
- Two ways of summarizing the ratios for the matched peers (taking either the mean or median).

In Figure 1 we report the valuation estimates for CVS using the 24 permutations of inputs. The top panel displays the estimates in order from smallest to largest, while the bottom panel records the specific combination of inputs for the estimates. In this example, experts could generate a range of estimates, from \$40 billion to \$107 billion, using input values consistent with the conventional textbook approach to valuation. The true valuation on date d^*, represented by the dashed line in the top panel, was \$48 billion.

Figure 1: **Input Choice and Valuation** This figure reports the valuation estimates using different input combinations for CVS seven trading days after the end of 2010 Q4. The top panel plots the estimated market value from each of the 24 combinations, which are represented by the grey tiles in the lower panel. The combinations vary based on the choice of matching feature set, number of matches, distance measure, and summary measure.

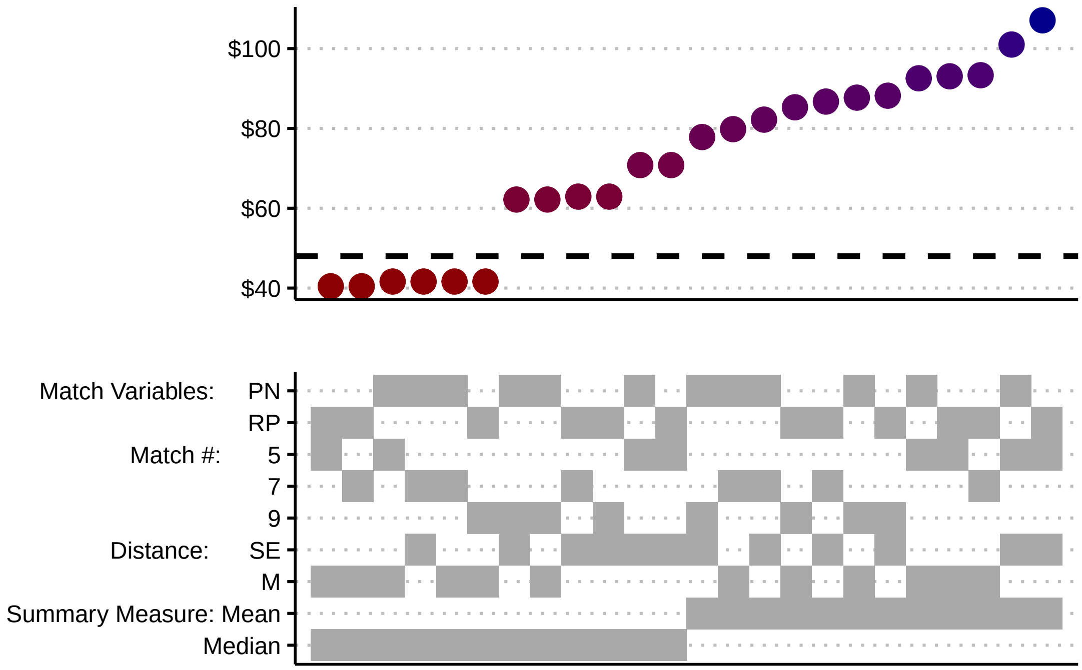

As evidenced by Figure 1, the conventional approach to valuation leaves substantial areas of expert discretion, which lead to large differences in valuation estimates. We believe that this fact explains why valuation reports in litigation typically vary so widely between defense and plaintiff experts; even cabining consideration to strategies squarely within textbook best practice, the range of potential estimates swamps any underlying signal.

We estimate a measure of the range of the valuation distribution for each of our 10,000 randomly-selected firm quarters to explore the general impact of this discretion. For each observation, we calculate the 24 unique valuation estimates using the permutations described earlier (two choices of matching variables, three choices of matching number k, two different distance measures, and two summary measures). We then assume that a defense expert picks the second-lowest estimate for their report, and the plaintiff chooses the second highest.[13]

Figure 2: **Separating Distributions** This figure reports the distribution of the second-lowest and second-highest valuation estimates from the permutations of the conventional k-NN approach, measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. We assume that the defense expert picks the second lowest valuation, and the plaintiff expert selects the second highest.

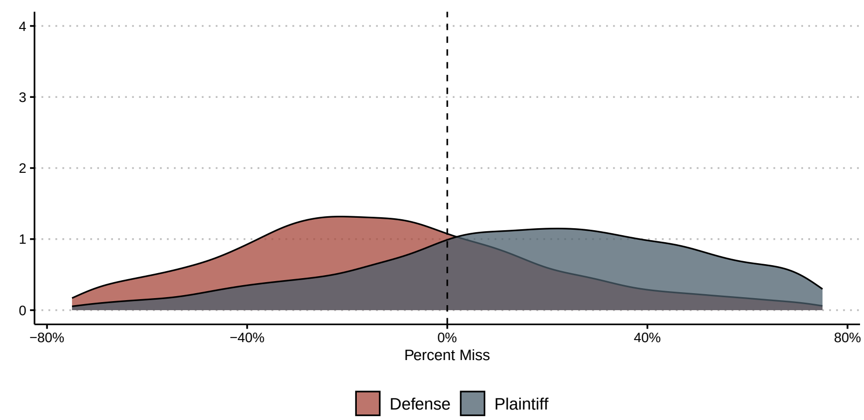

Kernel density estimates for the defense and plaintiff-friendly valuations in our Monte Carlo sample are reported in Figure 2. To allow for different scales of randomly-selected observations, we focus on percent deviations from the true market capitalization value on date d^*. In addition, to minimize the role out outliers, we restrict our attention to estimates that are within 75% of the true value. The estimates separate into two distinct distributions, centered generally at $\pm$ 25%. Thus, using our sample and straightforward applications of conventional valuation practice, we can generate large differences in reported results, both of which would be viable under the *Daubert* standard of admissibility. This suggests that the common criticism of valuation disparities by the judiciary[14] could simply be driven by experts exploiting the degrees of freedom afforded by the conventional k-NN approach.

One potential remedy to this problem could be to simply take an average of the estimates from the permutations. This would be analogous to a court-appointed

[13]An alternative way to interpret the strategy would be to assume that the litigation team surveys experts and selects the second lowest or highest estimate. Given the cost of producing expert analyses, this seems like the less plausible thought experiment.

[14]include citation

expert giving an impartial summary of the range of valid estimates. To explore this possibility, we take the median of the 24 estimates for each of our 10,000 randomly-selected observations, which we report in Figure 3. In Panel A we superimpose the density of the median of the estimates onto the curves for the plaintiff and defense friendly estimates. Over our sample of firms, the median of the permutations represents a plausibly unbiased estimate for the true value. However, while centered at the true value, there is a large variance around the mean, as demonstrated by the fat tails of the distribution. In Panel B we report that 11% of the estimates are more than 40% below the true value, while 24% are greater than 40% above.

Figure 3: **Averaging over Permutations** This figure reports the distribution of the median valuation estimate from the permutations of the conventional k-NN approach, measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. In Panel A we report this median estimate over the plaintiff and defense density estimates. In Panel B, we report the mass of the distribution falling outside of $\pm$ 40% of the true market capitalization.

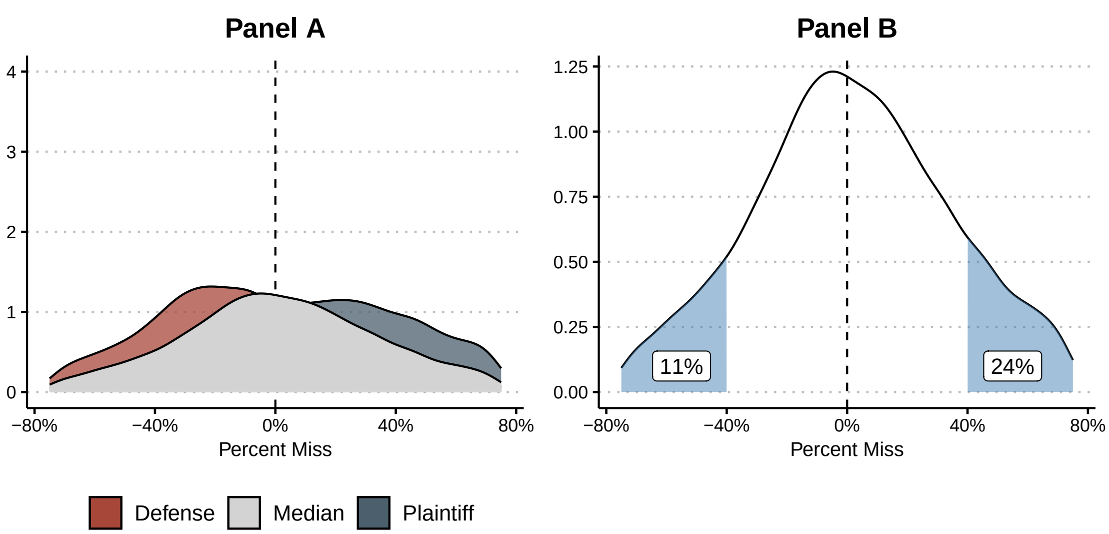

4.2 Data-Driven Approaches

As we saw in Section 2, the conventional approach to valuation suffers from a number of methodological challenges in the litigation context, namely that it provides large areas of discretion for experts, and produces estimates with large variance. In this section, we explore whether modern data-driven approaches to prediction can address some of these concerns.

4.2.1 Quarterly Data

We begin our extensions by building off the conventional approach that uses quarterly financial data from Compustat. Our first analytical point is to simply note that we actual have *time series* of valuation ratios. The conventional k-NN approach to valuation identifies peer firms using sets of covariates at one discrete point in time, generally the closest fiscal quarter to the valuation date. However, we have repeated measures of the valuation measures, which we can use to more accurately impute predictions in later periods.

In Figure 4 we plot the time series of the ratio of market capitalization to EBIT for CVS (in blue) as well as its industry peer firms for the eight fiscal quarters preceding the target data. The ratio hovers around 35 for CVS over this period, and a number of peers have ratios similar in magnitude to CVS, with a few peers exhibiting much larger ratios. This fact likely explains why the estimates from the conventional k-NN approach in Figure 1 exhibits such a large right tail (especially when focusing on the arithmetic mean of the ratio for the closest peer).

Figure 4: **Time Series Dynamics of Valuation Ratio** This figure reports the quarterly valuation ratio (market capitalization over EBIT) measures for CVS and its peer comparables. The time series for CVS is represented in blue while the peers are represented in gray. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter, and d^* represents the randomly chosen valuation date.

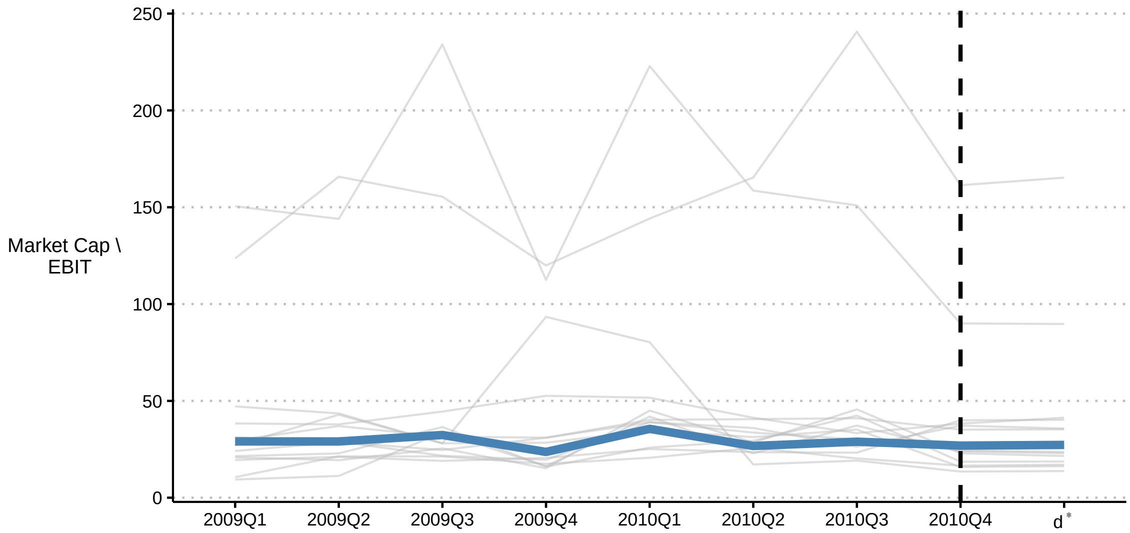

The time series properties exhibited by CVS and its peers in Figure 4 suggest a straightforward method for imputing the predicted ratio on the valuation date d^*; rather than using a set of (clearly arbitrary) covariates to generate a matching set, we simply use the valuation ratios for CVS and its peers in the spirit of a "synthetic controls" approach (CITE ABADIE, Imbens Doudchenko). More formally, we estimate penalized regression models of the form:

$$\widehat{\beta} = \text{argmin}_{\beta} \left(||y - X\beta||^2 + \lambda \left(\alpha ||\beta||_1 + \left(\frac{1-\alpha}{2} \right) ||\beta||^2 \right) \right) \tag{1}$$

where y is the 8×1 vector of valuation ratios for CVS, X is the $8 \times p$ matrix of quarterly valuation ratios for the peer firms (where p is the number of peer firms), and β is a vector of coefficient values on the peer firms.

The first term in the objective function is a minimization of the squared loss between the target valuation ratio and the prediction using a linear combination of the ratio of the peer firms. The second portion penalizes the coefficient values to prevent overfitting. Penalized regression is indeed necessary for the matrix to be invertible when $p > n$.

In Equation 1 the λ term controls the magnitude of the penalty, and the α term controls how the penalization occurs. Setting $\alpha = 1$ is lasso regression and the penalization is done through the L_1 norm, while setting $\alpha = 0$ is ridge regression which uses the L_2 norm. The use of intermediate values of $\alpha \in (0, 1)$ is generally referred to elastic net, and in our empirical results we choose an optimal value of α in the estimation period using a grid search over $[0, 1]$. In addition, we choose the optimal value of λ using leave-one-out cross-validation over the eight quarterly observations. In general, lasso regression tends to shrink the coefficient estimates in $\widehat{\beta}$ towards 0, while ridge will shrink them towards each other.

Table 4: Penalized Regression Weights on Peer Firms

Company	Lasso	Ridge	Elastic Net
INTERCEPT	27.7365	28.9667	27.8344
AMAZON.COM INC	0.0000	0.0000	0.0000
BIG 5 SPORTING GOODS CORP	0.0000	0.0000	0.0000
BLUE NILE INC	0.0082	0.0001	0.0076
CABELAS INC	0.0000	0.0001	0.0000
CASH AMERICA INTL INC	0.0000	0.0001	0.0000
FIRSTCASH HOLDINGS INC	0.0000	0.0002	0.0000
GLOBAL INDUSTRIAL CO	0.0000	0.0004	0.0000
HSN INC	0.0000	0.0002	0.0000
MEDCO HEALTH SOLUTIONS INC	0.0000	0.0001	0.0000
OMNICARE INC	0.0000	-0.0001	0.0000
PCM INC	0.0000	0.0000	0.0000
QURATE RETAIL INC	0.0000	0.0002	0.0000
TANDY LEATHER FACTORY INC	0.0000	0.0002	0.0000

This table reports the coefficient values on the peer firms using different forms of penalized regression. The outcome variable is the ratio of market capitalization to EBIT for the target firm, and the features that enter the regression are the ratios for the peer firms. We use quarterly data for the preceding two years in fitting the model, and optimize the tuning parameter using leave-one-out cross validation.

In Table 4 we report the coefficient values $\widehat{\beta}$ and model intercept for each of the three penalized-regression models for our motivating example. The lasso model shrinks most of the coefficient values to 0, except for the coefficient on the valuation ratio for Blue Nile Inc., an online jewelry retailer. Elastic net produces nearly identical results—the optimal value for α in this example is 0.9. As expected, the ridge regression model shrinks the coefficients towards each other, so there are fewer firms that drop out of the model entirely.

We report the time series values for the CVS valuation ratio (in red) as well as the model fitted estimates in Figure 5. Consistent with how penalized models trade-off bias for lower variance, we see that the valuation ratio estimates shrink towards the time-series average, while the underlying ratio exhibits larger swings. In general, to prevent overfitting the data when doing out-of-sample prediction, this is precisely what simple, well-calibrated models are designed to do. All of the models produce similar estimates, which end up being within 6% of the true valuation ratio on date d^*.

Figure 5: **Valuation Ratio and Penalized Regression Predictions** This figure reports the quarterly valuation ratio (market capitalization over EBIT) for CVS, as well as its predicted value using different penalized regression models. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter, and d^* represents the randomly chosen valuation date.

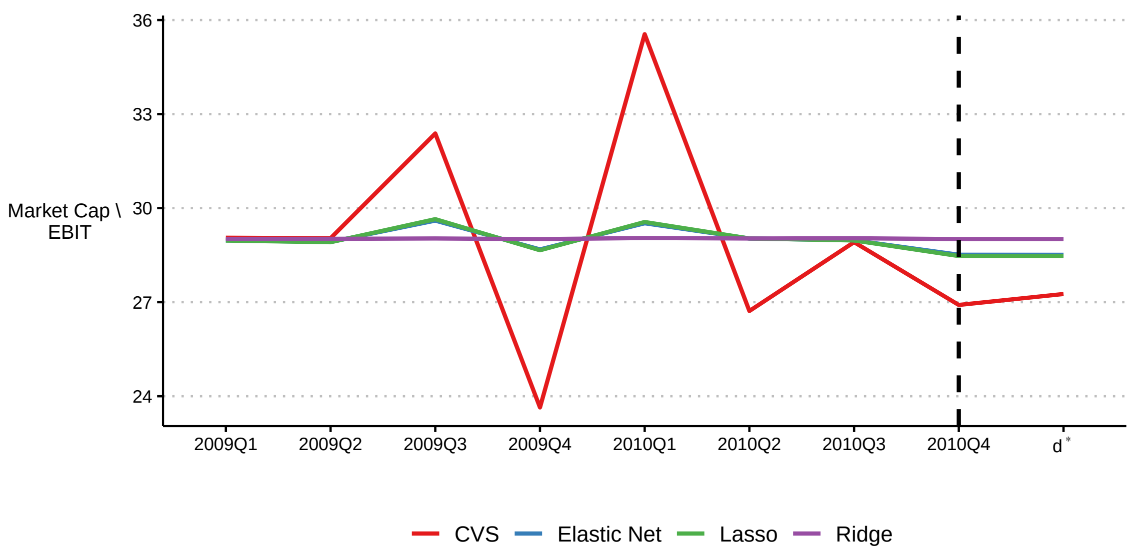

These estimates are clearly much closer to the true value in this example than the conventional k-NN estimates from Figure 1. However, it is not possible to determine whether the data-driven approach is superior based on a single example. When doing valuation, we want the best estimate as of the end of the preceding fiscal quarter. Intervening events in the data could make a specification appear superior even with a worse underlying prediction given the available data as of the prediction date. Our Monte Carlo analysis allows us to estimate similar examples over the 10,000 randomly-drawn observations, which will average away stochastic changes in the valuation that are incapable of accurate prediction.

In Figure 6, Panel A, we overlay the density estimate from the lasso regression predictions over the plaintiff, defense, and median curves generated previously. Even with only eight quarterly observations, the lasso model performs visibly better than the simple median of the 24 k-NN permutations, with a peak centered roughly at zero and much lower variance (i.e. thinner tails). In Panel B we produce the density estimates for the lasso, ridge, and elastic net models. Using quarterly data and targeting the valuation ratio, the elastic net model seems to perform slightly better than lasso, which performs better than ridge. All three approaches however outperform the conventional k-NN approach.

Figure 6: **Kernel Density Estimates Using Penalized Regression** This figure reports the kernel density estimates for the conventional approach as well as for the penalized regression models. The reported value is measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. In Panel A we report the lasso regression model estimates over the plaintiff, defense, and median density estimates. In Panel B, we report the three different penalized models separately.

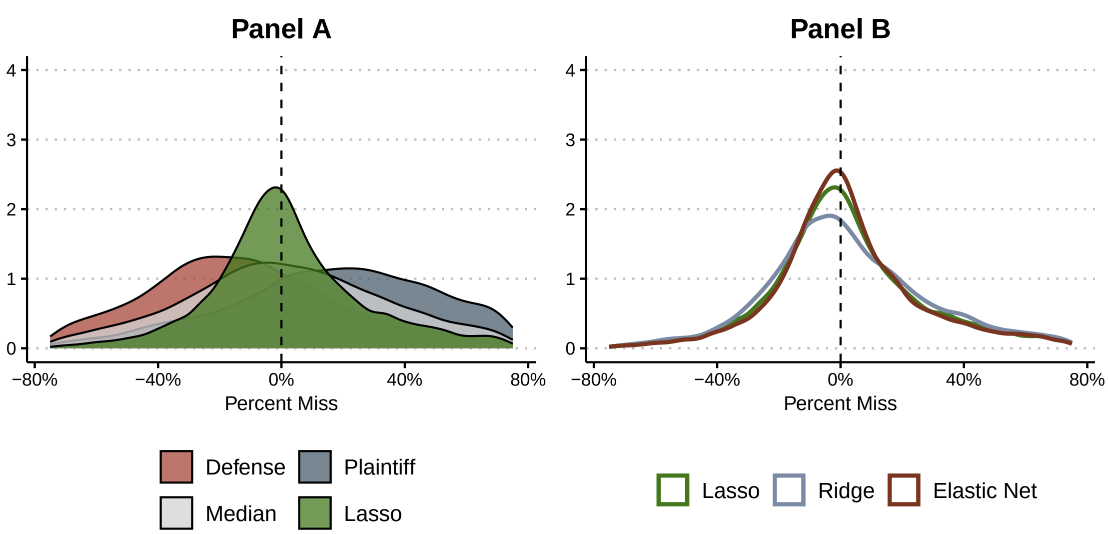

An obvious question is whether we can generate better estimates by avoiding the use of valuation ratios entirely. The goal of a valuation for litigation purposes is to generate a best-estimate for the market capitalization (or enterprise value) of a firm absent the intervention of some event. The use of ratio scaling is required under the conventional approach to allow for scale differences among firms of different sizes, as the conventional k-NN approach relies on a convex weighting scheme. However, the penalized regression approach does not suffer from the same limitations, as the intercept term (and the common use of standardization before the minimization) flexibly controls for such scale differences.

In Figure 7 we report the quarterly market capitalization value for CVS and its industry peers for the two-year period leading up to the valuation date. In comparison to the time-series trends in valuation ratio, the market capitalization is relatively stable. In addition, a number of peers are clearly much smaller (in market capitalization terms) than CVS, with a few peers of roughly similar size.

Figure 7: **Time Series Dynamics of Market Capitalization** This figure reports the quarterly market capitalization for CVS and its peer comparables. The time series for CVS is represented in blue while the peers are represented in gray. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter, and d^* represents the randomly chosen valuation date.

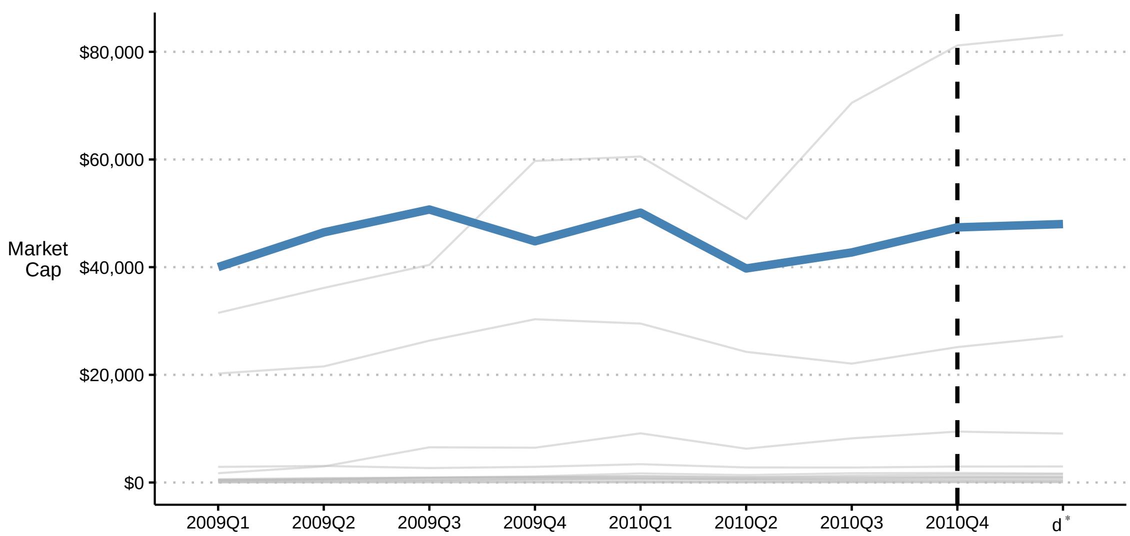

In Table 5 we report the coefficient estimates from the penalized regression models specified through Equation 1 using quarterly market capitalization as the outcome variable. In comparison to Table 4, more peer firms remain receive non-zero weight following penalization, suggesting that the predicted market capitalization values may not be as static as the ratio valuation model. In addition, the lasso and elastic net models collapse to the same estimate, because the optimal value α in this sample is 1 (or equivalent to lasso penalization).

Table 5: Penalized Regression Weights on Peer Firms (Market Cap)

Company	Lasso	Ridge	Elastic Net
INTERCEPT	19781.4724	36854.8411	19781.4724
AMAZON.COM INC	0.0000	0.0005	0.0000
BIG 5 SPORTING GOODS CORP	0.0000	2.5620	0.0000
BLUE NILE INC	6.7892	1.6834	6.7892
CABELAS INC	0.0000	0.3236	0.0000
CASH AMERICA INTL INC	0.0000	0.4500	0.0000
FIRSTCASH HOLDINGS INC	-9.7636	-0.2475	-9.7636
GLOBAL INDUSTRIAL CO	19.1242	0.8182	19.1242
HSN INC	0.0000	0.0582	0.0000
MEDCO HEALTH SOLUTIONS INC	0.0000	0.0566	0.0000
OMNICARE INC	0.0000	0.8017	0.0000
PCM INC	228.8674	20.0044	228.8674
QURATE RETAIL INC	0.0802	0.0527	0.0802
TANDY LEATHER FACTORY INC	0.0000	-5.5843	0.0000

This table reports the coefficient values on the peer firms using different forms of penalized regression. The outcome variable is market capitalization for the target firm, and the features that enter the regression are the maket capitalizations of the peer firms. We use quarterly data for the preceding two years in fitting the model, and optimize the tuning parameter using leave-one-out cross validation.

Figure 8 plots the quarterly market capitalization for CVS (in red) along with the imputed market capitalization estimates from the three penalized regression models.[15] Here the ridge regression estimates (in purple) still shrink towards the sample average, while the the lasso and elastic net models produce estimates that closely match the observed valuation over the sample period. In addition, the estimates on the valuation date are much closer for the lasso and elastic net estimates.

[15] Note that because the elastic net model produces identical estimates to the lasso model, their lines are coextensive.

Figure 8: **Market Capitalization and Penalized Regression Predictions** This figure reports the market capitalization for CVS, as well as its predicted value using different penalized regression models. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter, and d^* represents the randomly chosen valuation date.

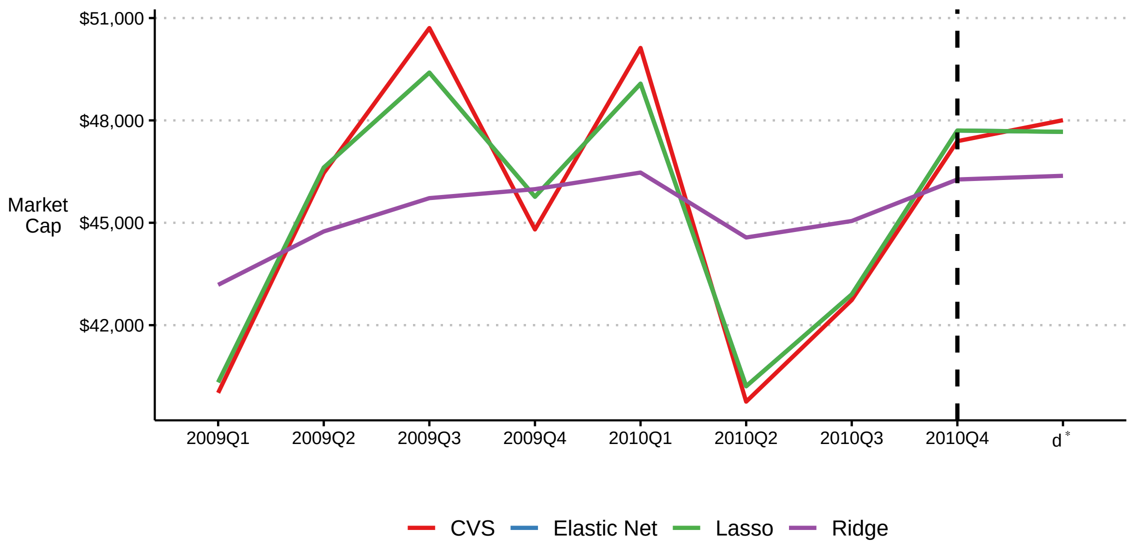

To explore whether these results hold more generally, in Figure 9 we report the density estimates for the models with market capitalization as the outcome variable. Panel A provides the density estimates for the lasso regression model with market capitalization as the outcome over the estimates from Figure 5. Over the Monte Carlo sample, directly targeting firm value produces estimates with substantially lower variance and a higher peak around the true valuation levels. In Panel B we report each of the penalized market capitalization models separately. Over the full sample the elastic net model (i.e. choosing an optimal α value to minimize leave-one-out prediction error) produces the best estimates, followed by lasso and ridge.

Figure 9: **Kernel Density Estimates Using Penalized Regression** This figure reports the kernel density estimates for the conventional approach as well as for the penalized regression models. The reported value is measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. In Panel A we report the lasso regression model estimates using both the ratio and market capitalization as the outcome variable over the plaintiff, defense, and median density estimates. In Panel B, we report the three different penalized models with market capitalization as the outcome measure separately.

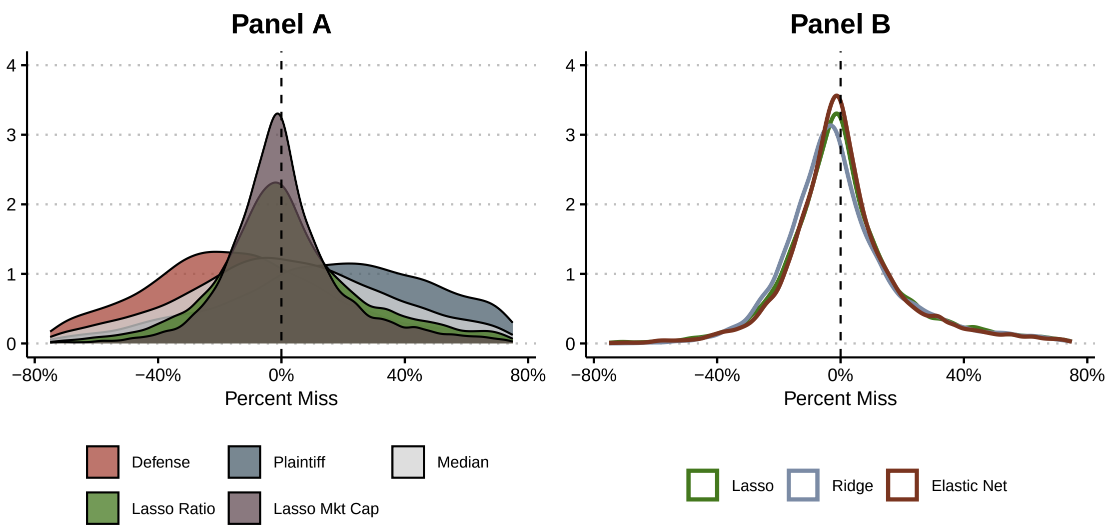

4.2.2 Daily Data

As shown in Figure 9, using market capitalization rather than the valuation ratio as the outcome variable leads to substantial gains in accuracy for predicting firm value. Given that this approach avoids the use of quarterly financial reporting data entirely, we can potentially achieve more accurate results by using *daily* stock price data, rather than quarterly measures. In particular, machine learning algorithms tend to perform better with more granular data, suggesting that the move to higher frequency measurements can further increase predictive performance.

In Figure 10 we plot the daily market capitalization values for CVS and its industry peers for the one-year period predceding the randomly-chosen fiscal year end up to the valuation date. Figure 10 is simply the daily-measured corollary to Figure 7, and we again see that CVS is larger than most of its peers over the relevant time period, with a comparatively stable valuation path.

Figure 10: **Daily Time Series Dynamics of Market Capitalization** This figure reports the daily market capitalization for CVS and its peer comparables. The time series for CVS is represented in blue while the peers are represented in gray. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter.

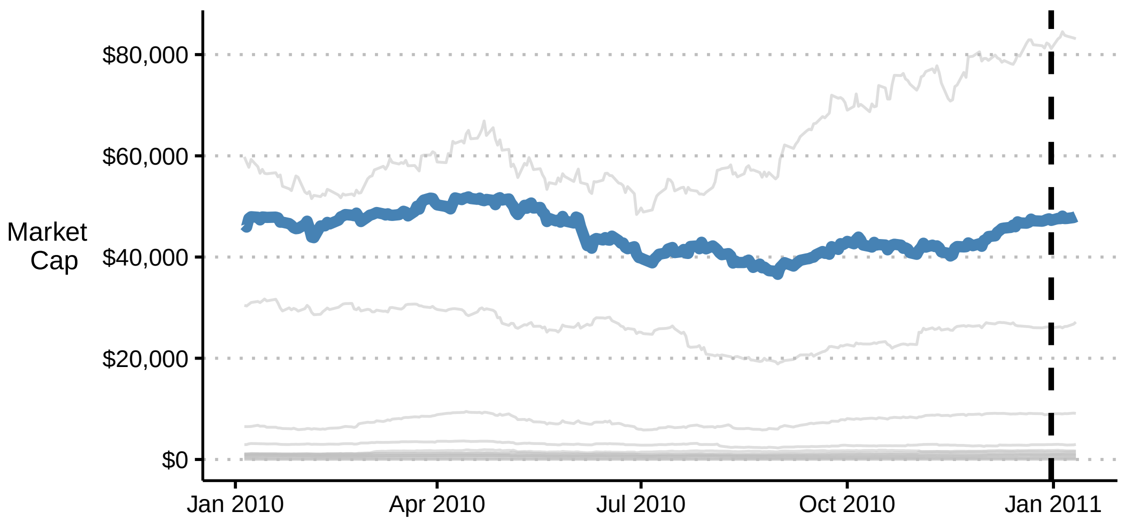

Table 6 provides the model intercept and coefficient estimates on the industry peers for our motivating example using the penalized regression models. The outcome variable is now the daily market capitalization value for CVS, and the estimation period spans the 250-day trading period ending on the fiscal quarter end. Given the longer sample period in these regressions, we use 25-fold cross-validation, rather than more time-intensive leave-one-out optimization.[16] The models are now much less sparse, with each industry peer getting positive (in absolute value terms) weight following the penalization. In addition, the estimates appear largely correlated with each other, with the elastic net model again being closer to lasso than ridge.

[16]In unreported results we test whether time-series cross-validation models generate superior estimates to simple cross-validation. While these models produce slightly lower-variance estimates, the differences are marginal.

Table 6: Penalized Regression Weights on Peer Firms (Daily Market Cap)

Company	Lasso	Ridge	Elastic Net
INTERCEPT	10458.2596	11506.3925	10553.5602
AMAZON.COM INC	0.1189	0.0176	0.1140
BIG 5 SPORTING GOODS CORP	40.2609	36.8099	40.1689
BLUE NILE INC	4.5676	5.7553	4.6317
CABELAS INC	2.7735	1.1663	2.7146
CASH AMERICA INTL INC	4.2031	3.4243	4.0575
FIRSTCASH HOLDINGS INC	-15.7395	-7.3636	-15.4644
GLOBAL INDUSTRIAL CO	8.4333	6.7069	8.3250
HSN INC	2.1101	0.0179	2.0337
MEDCO HEALTH SOLUTIONS INC	0.1873	0.1592	0.1866
OMNICARE INC	0.8669	1.1434	0.8910
PCM INC	61.0722	37.0093	60.7036
QURATE RETAIL INC	-0.9016	0.0850	-0.8570
TANDY LEATHER FACTORY INC	43.3947	52.6340	44.2837

This table reports the coefficient values on the peer firms using different forms of penalized regression. The outcome variable is daily market capitalization for the target firm, and the features that enter the regression are the maket capitalizations of the peer firms. We use daily data starting 250 trading days before fiscal quarter end in fitting the model, and optimize the tuning parameter using 25-fold cross validation.

Figure 11 uses the fitted coefficient estimates from Table 6 to impute the expected market capitalization for CVS over the estimation and valuation period. In addition, we include the predictions using a non-regression based machine learning estimate (a random forest model)[CITE]. As evidenced by the plot, the models are largely capable of predicting the time series properties of CVS's valuation (in red) over the relevant time period. In particular, the random forest model is able to very closely match the daily valuation of CVS over the estimation period, perhaps at the risk of overfitting the data.

Figure 11: **Daily Market Capitalization and Model Predictions** This figure reports the daily market capitalization for CVS, as well as its predicted value using different penalized regression models. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter, and d^* represents the randomly chosen valuation date.

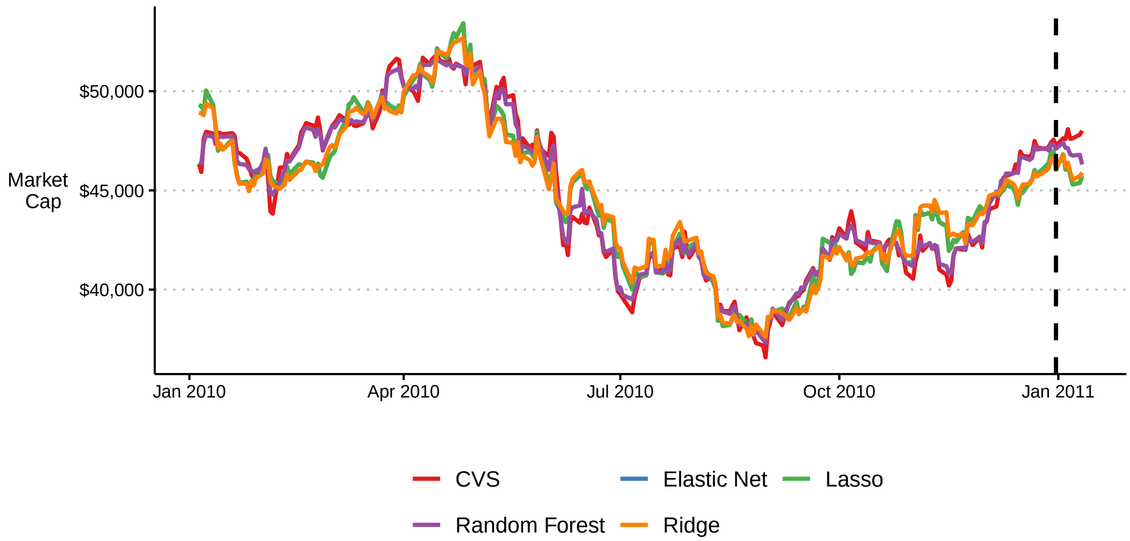

We again might ask how well the results from the motivating example generalize to the typical case. In Figure 12, we report the kernel density estimates for the valuation predictions from the conventional approach and the daily market capitalization models. In Panel A we report the conventional k-NN estimates, as well as the density for the lasso regression model using both quarterly and daily values of firm market capitalization. The use of daily data further improves on the performance of the data-driven approach to valuation, as the daily lasso regression estimates are higher peaked around zero than the quarterly measures. In Panel B we report separately the density for each model estimate using daily market capitalization as the outcome variable. The random forest and ridge regression models have the best, roughly equivalent, performance over our Monte Carlo sample.

Figure 12: **Kernel Density Estimates Using Daily Data Penalized Regression** This figure reports the kernel density estimates for the conventional approach as well as for the penalized regression models. The reported value is measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. In Panel A we report the lasso regression model estimates using both daily and quarterly market capitalization as the outcome variable over the plaintiff, defense, and median density estimates. In Panel B, we report the three different penalized models, as well as the random forest estimates, with market capitalization as the outcome measure.

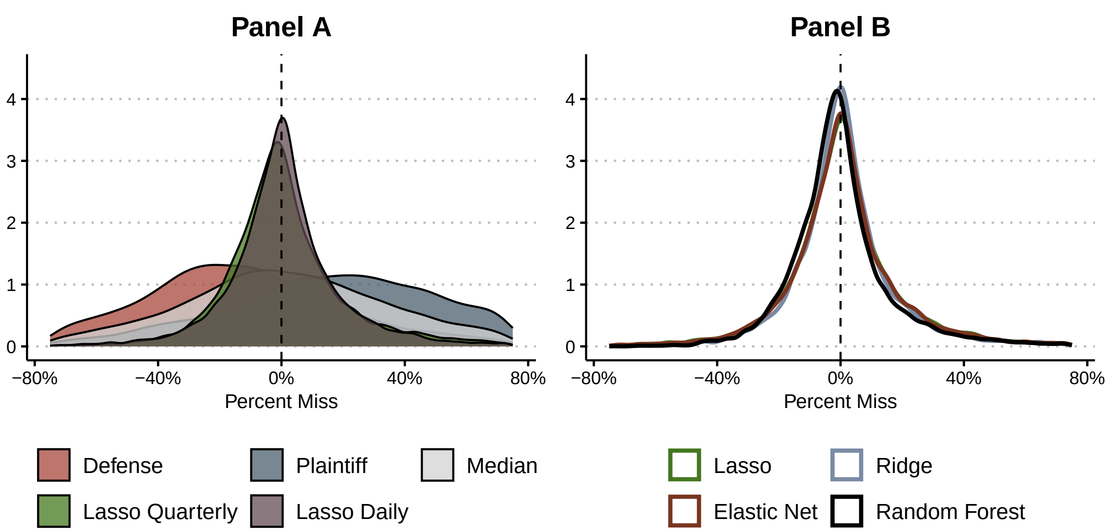

Given that we are now using daily pricing data, it is not clear why we should continue to use the valuation *level*, rather than modeling the relationship between firm *returns* and market and peer factors. The finance literature almost exclusively focuses on returns, given the long right tail of financial data (CITE). In fact, the other primary area of litigation that uses financial prices to measure damages—securities litigation—has consistently used returns-based modeling to capture the value relevance of disputed events. There is no obvious ex-ante rationale for why litigation surrounding disputed firm valuations should differ methodologically.

To explore the advantage of using daily returns in the valuation context, we supplement our Monte Carlo analysis with predictions based on imputing counterfactual return series from the model-implied factor loadings of returns-based regressions of the Fama-French-Carhart variety. Assume that r_{it} represents the return on firm i's stock on date t. Assume also that the target firm is denoted by J. In our regressions we include the market return (MKTRF), as well as the factor returns from (CITE FAMA-FRENCH) and (CITE CARHART), which include returns on long-short portfolios sorted by size (SMB), value (HML), and momentum (UMD). Finally, as shown in (CITE BAKER AND GELBACH), the returns on peer firms (either an equally-weighted average of the peer firms returns, or controlling for the peers separately) increases the predictive power of event study models.

We test five different models using returns as the outcome variables:

1. FFC Index:

   $$y_{i=J,t} = \alpha + \beta_1 MKTRF_t + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 UMD_t + \beta_5 PEERINDEX_t + \epsilon_{it}$$

   where the peer index is the daily equally-weighted average of the returns on the target firm's industry peers.

2. FFC All Peers:

   $$y_{i=J,t} = \alpha + \beta_1 MKTRF_t + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 UMD_t + \Omega r_{i \neq J,t} + \epsilon_{it}$$

   where the daily returns for the target firm's industry peers enter the regression separately.

3. Lasso:

   $$\text{argmin}_\beta \left(||r_{i=J,t} - X\beta||^2 + \lambda ||\beta||_1 \right)$$

   where the factor matrix X is the union of an intercept, the Fama-French-Carhart factors and the individual peer firm returns. The penalization is done through the L_1 norm.

4. Ridge:

   $$\text{argmin}_\beta \left(||y - X\beta||^2 + \lambda \left(\frac{1-\alpha}{2} \right) ||\beta||^2 \right)$$

   where the factor matrix X is the union of an intercept, the Fama-French-Carhart factors and the individual peer firm returns. The penalization is done through the L_2 norm.

5. Elastic Net:

   $$\text{argmin}_\beta \left(||y - X\beta||^2 + \lambda \left(\alpha ||\beta||_1 + \left(\frac{1-\alpha}{2} \right) ||\beta||^2 \right) \right)$$

where the factor matrix X is the union of an intercept, the Fama-French-Carhart factors and the individual peer firm returns. The penalization is a convex average of the L_1 and L_2 norm, where α is chosen optimally based on cross-validation error in the estimation sample period.

The first two models are estimated using simple ordinary least squares, while the lasso, ridge, and elastic net models are penalized regressions following Equation 1. Table 7 reports the coefficient estimates from the five returns-based models for our motivating example. In comparison to the daily market capitalization results in Table 6, the feature matrix appears much more sparse, as the lasso and elastic net model both put zero weight on the portfolio factor returns and a number of the industry peers. Again, the optimal-α elastic net model is much closer to lasso than ridge.

Table 7: Return Coefficients

Company	FFC Index	FFC All Peers	Lasso	Ridge	Elastic Net
INTERCEPT	-0.0003	-0.0003	-0.0003	-0.0003	-0.0003
MKTRF	0.8042	0.9051	0.7055	0.6452	0.6796
SMB	-0.0303	0.0149	0.0000	-0.0128	0.0000
HML	-0.1407	-0.0992	0.0000	0.0038	0.0000
UMD	-0.1637	-0.1843	0.0000	-0.0699	0.0000
PEER INDEX	0.1811				
AMAZON.COM INC		0.0453	0.0151	0.0559	0.0203
BIG 5 SPORTING GOODS CORP		-0.0084	0.0000	0.0003	0.0000
BLUE NILE INC		0.0039	0.0000	0.0119	0.0000
CABELAS INC		0.0655	0.0430	0.0611	0.0443
CASH AMERICA INTL INC		-0.0107	0.0000	0.0030	0.0000
FIRSTCASH HOLDINGS INC		-0.0493	0.0000	-0.0219	0.0000
GLOBAL INDUSTRIAL CO		0.0126	0.0050	0.0265	0.0084
HSN INC		0.0033	0.0000	0.0066	0.0000
MEDCO HEALTH SOLUTIONS INC		-0.0245	0.0000	-0.0040	0.0000
OMNICARE INC		0.0668	0.0408	0.0747	0.0446
PCM INC		0.0256	0.0135	0.0263	0.0140
QURATE RETAIL INC		-0.0755	0.0000	-0.0460	0.0000
TANDY LEATHER FACTORY INC		0.0393	0.0245	0.0393	0.0247

This table reports the coefficient values on the Fama-French-Carhart factors, and the peer firms, using both ordinary least squares and different forms of penalized regression. The outcome variable is the return on CVS's stock, and the features that enter the regression are the Fama-French-Carhart factors and the returns for the peer firms. We use daily data for the 250 days prior to and ending on the fiscal-year end in fitting the model. For the penalized regression models we optimize the tuning parameter using 25-fold cross validation.

Using the factor loadings in Table 7, we can easily create predictions for the valuation on the target date. For each model, the coefficient estimates generate predictions for the return on the target firm's stock in the post-estimation period, $\widehat{r_{i=J,t}}$. Using the predicted returns, we calculated cumulative predicted returns for date d^* as:

$$\widehat{cr_{i=J,t=d^*}} = \prod_{t=\underline{\mathrm{t}}}^{d^*}(1 - \widehat{r_{i=J,t}} - div_{i=J,t})$$

where $\underline{\mathrm{t}}$ is the first trading date after the estimation period (i.e. the first trading date after the randomly-chosen fiscal quarter end date), d^* is the valuation date, and $div_{i=J,t}$ are dividends paid per share by the target firm on date t. We then simply multiply $\widehat{cr_{i=J,t=d^*}}$ by the trading price on the fiscal end date (making sure to adjust for intervening stock splits and issuances) to arrive at an estimated market capitalization for the target firm on the valuation date. We report the model predictions (colored lines) and the true market capitalization (black line) for CVS in the post-estimation period in Figure 13. The models generally, though not perfectly, capture the trend in the underlying valuation of CVS over this time period.

Figure 13: **Daily Market Capitalization and Returns Model Predictions** This figure reports the daily market capitalization for CVS and the prediction from a series of models using returns rather than valuation as the outcome variable. The time series for CVS is represented in black while the model predictions are in colors. The vertical dashed line represents the end of quarter-end for the randomly sampled quarter.

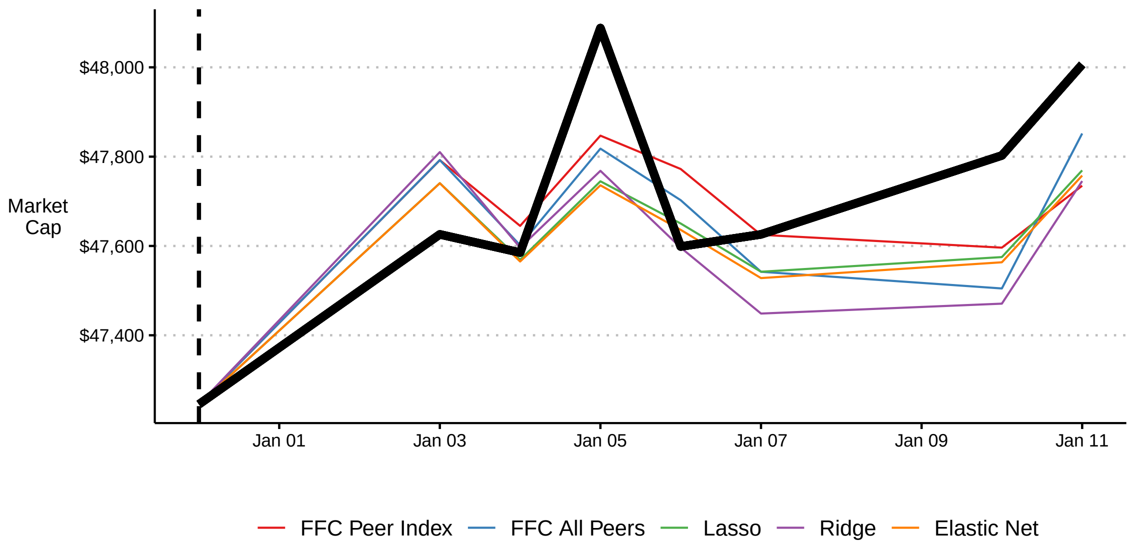

We report the kernel density estimates for the deviations of the predicted market capitalization values for our 10,000 randomly selected observations in Figure 14. Panel A reports the density for the conventional k-NN estimates, as well as for the lasso model with daily market capitalization as the outcome variables, and with daily returns as the outcome variable. The estimates that use daily returns instead of daily market capitalization are clearly superior (i.e. they are higher-peaked at zero). These results suggest that we may want to unify the approaches of valuing events and firms, at least for purposes of litigation.

Figure 14: **Kernel Density Estimates Using Daily Data** This figure reports the kernel density estimates for the conventional approach as well as for the returns-based models. The reported value is measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. In Panel A we report the lasso regression model estimates using both daily market capitalization and daily returns as the outcome variable over the plaintiff, defense, and median density estimates. In Panel B, we report the different returns-based model estimates separately.

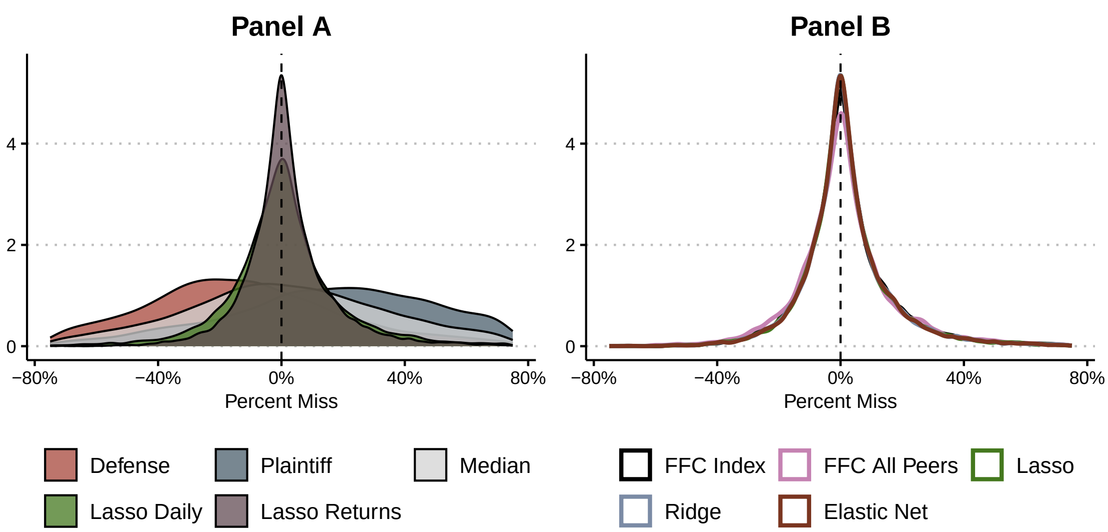

5 Real-World Application: *DFC Global*

Although our analysis offers a general approach to market-based valuation, the techniques illustrated above can be deployed helpfully in real-world situations that are far more targeted in nature. We explore one such application below, revisiting the landmark Delaware case of DFC Global Corp. v. Muirfield Value Partners, L.P.,[17] a shareholder dispute that metastasized into a famously focal flashpoint for valuation methodology. *DFC* marks something of a watershed moment in stockholder appraisal cases—statutorily authorized actions that are brought by dissenting stock-

[17] 172 A.3d 346 (Del. 2017).

holders after the close of certain eligible transactions.[18] By statutory command, the appraisal inquiry focuses on assessing the "fair value" of the target as a going concern, using "all relevant factors," and specifically excluding the value of merger synergies or takeover premiums.[19] Accordingly, comparable companies analysis is an oft-utilized tool for reckoning valuation in appraisal proceedings. Both opposing experts in *DFC* utilized CC as part of their valuation analyses, leaving the court to grapple with their divergent opinions.[20]

5.1 Background

DFC Global,[21] a publicly traded payday lending firm, faced significant headwinds in 2012-13, including issues related to its financial leverage and regulatory scrutiny in several countries. In response, DFC engaged financial advisor Houlihan-Lockey to advise on a potential sale and initiate a bidding process. The bidding process was tumultuous, buffeted by several negative shocks and disappointing earnings reports, which impelled several bidders to withdraw. After contacting over 45 potential buyers, Houlihan eventually corralled 2-3 serious contenders, including a private equity company named Lone Star Funds ("Lone Star"). Ultimately, DFC signed a cash deal with Lone Star at \$9.50 per share on April 1, 2014, closing on June 13, 2014. Several DFC stockholders perfected their appraisal rights (led by hedge fund Muirfield Value Partners), and the case landed in front of Chancellor Bouchard of the Delaware Court of Chancery to determine fair value as of the closing date.[22]

5.2 Expert Opinions

Consistent with longstanding patterns in appraisal litigation under DGCL § 262, much of the substantive analysis in *DFC Global* came down to a valuation dance-off between opposing experts. Kevin Dages of Compass Lexecon, the petitioner's expert, performed a Discounted Cash Flow (DCF) valuation, estimating the fair value of DFC at \$17.90 per share. He also conducted a Comparable Companies analysis, choosing to peg DFC's EBITDA multiple at the 75th percentile among a set of 10 comparable companies that he had identified. Dages' CC analysis rendered a

[18]See Choi & Talley (2018).

[19]DGCL § 262

[20]It merits observing that CC and DCF approaches both continue to be part of the standard valuation canon in appraisal proceedings. *See, e.g,* HBK Master Fund v. Pivotal, C.A. No. 2020-0165-KSJM (Del. Ch. 2023) (ascribing equal weight to DCF and CC analyses).

[21]Ticker: DLLR; CIK: 0001271625; PERMNO: 1627099; PERMCO 46104

[22]In Re Appraisal of DFC Global, 2016 WL 3753123, Jul 08, 2016, at 12.

capacious value estimation range, comprising the interval between \$11.38 and \$26.95. Ultimately, Dages relied entirely on his DCF estimate, giving no weight in his final opinion to the CC method (even though a detailed CC analysis was included in his report). In rationalizing this decision, Dages asserted that "[t]he reliability of a multiples-based valuation is highly dependent on the ability to identify sufficiently comparable companies and transactions, or to properly adjust financial performance data to remove non-comparable items."[23]

Daniel Beaulne from Duff & Phelps provided the respondent DFC's expert report. Like Dages, he conducted booth a DCF and a CC analysis, which ultimately yielded estimates of \$7.81 per share and \$8.07 per share, respectively. Unlike Dages, Beaulne provided exclusively point estimates for his valuation approaches,[24] and he accorded equal weight to each of the DCF and CC estimates, ultimately delivering a fair value opinion of \$7.94 per share. For those keeping score at home, there was a cavernous gulf between Dages' and Beaulne's bottom-line fairness opinions—manifested in a valuation ratio of more than 2.25 to 1—a gap that is simultaneously gargantuan and entirely unremarkable in modern valuation cases.

5.3 Court of Chancery Opinion

Chancellor Bouchard delivered a 68-page opinion in July 2016 (two years after the deal closed). The Chancellor analyzed both experts' DCF and CC analyses, as well as the deal price itself, ultimately drawing from all three channels. As to CC, Bouchard sided with Beaulne's \$8.07 figure, largely because Dages was not able to justify his 75th percentile assumption (which he had never deployed in prior valuation reports[25]). In contrast, Bouchard substantially embraced Dages' DCF analysis, adapting it somewhat to deliver his own DCF estimate of \$13.07 per share. Chancellor Bouchard's opinion ultimately accorded equal weights to each of the three valuation lenses, with one-third weight apiece: DCF (\$13.07/share), Comparable Companies (\$8.07/share), and Deal Price (\$9.50/share). The end result was a blended \$10.21 per share assessment, handing a modest victory to the petitioners (which became less modest once augmented by statutory interest[26]). Following Lone Star's post-hearing

[23] Dages Report at 68.

[24] An analysis of Beaulne's expert report suggests that this decision was based in part on the fact that at least one of his earnings multiples rendered a *negative* equity value; rather than excluding it (which would have pushed the valuation higher), he instead averaged this negative multiple with the others. See Beaulne Report at 67-68

[25] Bouchard opinion, 2016 WL 3753123, at 56-57

[26] Under DGCL § 268(h), prejudgment interest in an appraisal action is compounded quarterly at the Federal Funds rate plus 500bps—a generous compounding factor for 2016, when spreads were

motion, which pointed out an error in the Court's DCF working capital projections, Chancellor Bouchard corrected his math in a revised opinion, but he simultaneously adjusted the perpetuity growth rate assumption as well, resulting in a post-correction valuation that came in at virtually the same figure as the original.[27]

5.4 Delaware Supreme Court's Decision

DFC appealed, making several forceful arguments. The most attention-grabbing of them was the contention that the Chancery Court's discretion should be limited in any appraisal-eligible acquisition that features an arm's-length sale following a competitive bidding process. In such situations, the appellants argued, the transaction price (less synergies) should be the definitive measure of "fair value" under the appraisal statute. This issue alone elicited significant attention, including dueling amicus briefs—one submitted by several law professors, and another submitted by a combined group of legal, economics and finance scholars (even including a Nobel laureate).[28] In the end, the Supreme Court substantially rejected DFC's categorical argument, emphasizing the criticality of preserving the Chancery Court's discretion.[29] At the same time, however, the Court admonished the Chancery Court to do a better job of "showing its work" to justify how it weighs different valuation approaches; moreover, it suggested that a competitive bidding process could well provide a sound reason to accord greater weight to the deal price than it might otherwise garner.[30] On the issue of post-hoc adjustments in the DCF perpetuity growth rate, the Supreme Court held that such changes were not supported by the factual record. The Court remanded the case back to Chancellor Bouchard for reconsideration.[31]

Key Takeaway: The Supreme Court's *DFC* opinion clearly reinforced the importance of discretion and the need for a detailed explanation by the fact finder of the preferred valuation method, clearly signaling that in some cases, a single metric may be most reliable, while in others multiple measuring perspectives should be consid-

---

extremely tight. See Jetley & Ji (2016)

[27]The revised opinion was actually $0.09 per share *higher* than the original, or $10.30. See 172 A.3d at 362.

[28]See Reynolds Holding, DFC Global Appraisal Battle Draws Opposing Briefs From Professors (Columbia Blue Sky Blog, Feb. 7, 2017).

[29]DFC Global Corporation v. Muirfield Value Partners, L.P. 172 A.3d 346 (Del. 2017). The Court noted that precisely the same argument had been made (unsuccessfully) to Chancellor Bouchard, and that the Supreme Court had at least twice rejected the proposition that merger price should categorically be conclusive in such situations. Id. at 364-65.

[30]For an analysis of this set of trade-offs in a theoretical setting, see Choi & Talley (2018)

[31]The case evidently settled shortly thereafter on undisclosed terms, upon which Chancellor Bouchard dismissed the action with prejudice.

ered. Perhaps more than anything, however, the case underscores the complexity of corporate valuation, and it highlights the necessity for judges to employ a thorough, principled approach in determining fair value, particularly when confronting gaping valuation chasms between competing expert opinions.

5.5 A *DFC* Do-Over?

How would the facts of *DFC Global* come out under our approach to CC valuation? Can one close the gulf between the expert valuation opinions? To investigate this claim, we use our data-driven approach to provide a valuation estimate for DFC Global at the valuation date. For our analysis we impute the DFC market capitalization using realized prices and returns for DFC and its peer firms, which are the union of the potential peer firms identified in the Dages and Beaulne reports[32] and the other financial firms in the same three-digit SIC code industry.[33]

First, we might ask what result would be achieved by using our penalized regression approach to valuation using quarterly data. In Figure 15 we report the quarterly market capitalization for the eight quarters ending on December 31, 2013, which we use as the last "unaffected" market price preceding the appraisal action. For visualization purposes, we report the market capitalization values as percent deviations from the starting period value (April 1, 2012). As is evident in the data, DFC substantially under-performed most, but not all, of its industry peers over this period.

[32]Due to data limitations, we omit non-US firms from the analysis. These firms were Cash Converters International Limited, International Personal Finance plc, and Provident Financial plc.

[33]The firm's three digit SIC industry as of fiscal-year 2013 was 609–Functions Related to Depository Banking. We chose not to use the two-digit SIC industry definition, because 60 (Depository Institutions) covered over 500 unique firms at that time.

Figure 15: **Quarterly Market Capitalization Over Time for DFC and Peers** This figure reports the percentage change in market capitalization for DFC and the identified peer firms for the two year period prior to calendar year end 2013, through the valuation date. The data is measured quarterly.

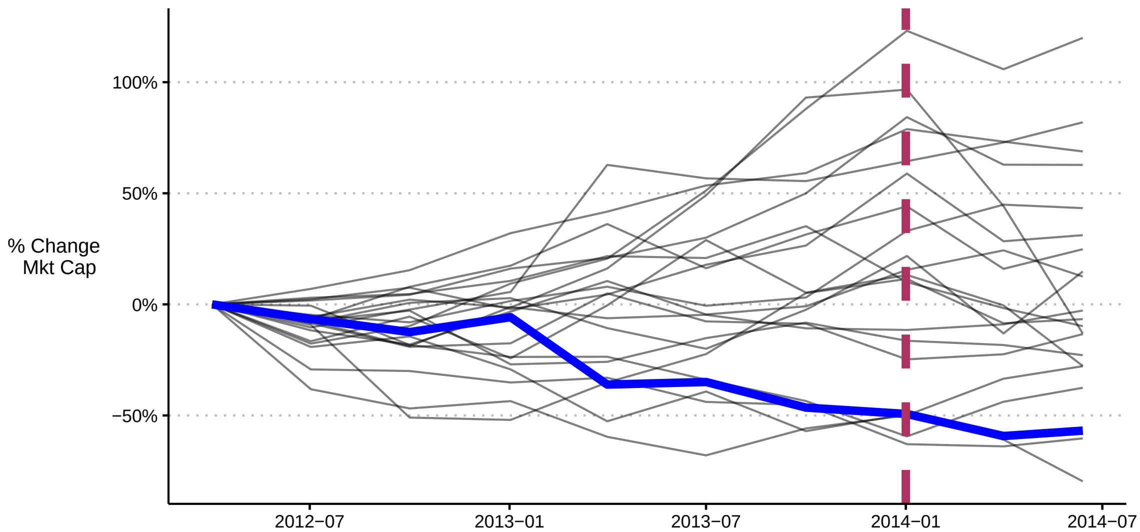

Similar to the analysis in Section 4.2.1, we can estimate the market capitalization for DFC on the target date (June 13, 2014) using the coefficient values from a regression of DFC's market capitalization on its industry peers using the quarterly market capitalization values. In Table 8 we report the fitted coefficients from the lasso, ridge, and elastic net models on the eight quarterly observations using leave-one-out cross-validation to optimize the penalty terms. The lasso model is sparse in this example, with only a handful of peer firms receiving positive weight following the penalization (as one would expect given the limited number of observations).

Table 8: Penalized Regression Weights on Peer Firms

Company	Lasso	Ridge	Elastic Net
INTERCEPT	894.9042	901.8631	1025.7439
ALBEMARLE CORP	0.0086	0.0104	0.0451
BLOCK H & R INC	-0.0402	-0.0051	-0.0157
CASH AMERICA INTL INC	0.0000	-0.0025	-0.0311
CREDIT ACCEPTANCE CORP	0.0000	-0.0165	-0.0177
EURONET WORLDWIDE INC	0.0000	-0.0128	-0.0150
EZCORP INC -CL A	0.0089	0.0221	0.0190
FIRST CASH FINANCIAL SVCS	0.0000	-0.0319	-0.0634
GLOBAL CASH ACCESS HOLDINGS	0.0000	0.0027	0.0000
GLOBAL PAYMENTS INC	0.0000	-0.0100	-0.0019
GREEN DOT CORP	0.0000	-0.0173	-0.0393
HEARTLAND PAYMENT SYSTEMS	0.0000	-0.0263	-0.0447
HIGHER ONE HOLDINGS INC	0.1618	0.0453	0.1283
MASTERCARD INC	0.0000	-0.0004	-0.0005
MONEYGRAM INTERNATIONAL INC	0.0000	-0.0359	-0.0921
QC HOLDINGS INC	0.0000	0.4664	0.2871
REGIONAL MANAGEMENT CORP	-0.5108	-0.0822	-0.1633
VERIFONE SYSTEMS INC	0.0000	0.0054	0.0116
VISA INC	0.0000	-0.0004	-0.0005
WESTERN UNION CO	0.0000	0.0016	0.0000
WORLD ACCEPTANCE CORP/DE	0.0000	-0.0705	-0.1627

This table reports the coefficient values on the peer firms using different forms of penalized regression. The outcome variable is the market capitalization for the target firm, and the features that enter the regression are the market capitalization levels for the peer firms. We use quarterly data for the preceding eight calendar quarters in fitting the model, and optimize the tuning parameter using leave-one-out cross validation.

In figure 16 we use the fitted coefficients from Table 8 to predict the market capitalization for DFC over the estimation period and through the valuation date. We see that the penalized models (in particular the lasso and elastic net models) do a generally good job at tracking DFC's valuation over this period. All three models report a valuation above the true realized valuation for DFC (the red line) on the valuation date. The quarterly models imply a valuation range of \$13.59 to \$15.08 a share.

Figure 16: **Quarterly DFC Market Capitalization and Model Predictions** This figure reports the quarterly market capitalization for DFC, as well as its predicted value using different penalized regression models. The vertical dashed line represents the end of calendar year 2013.

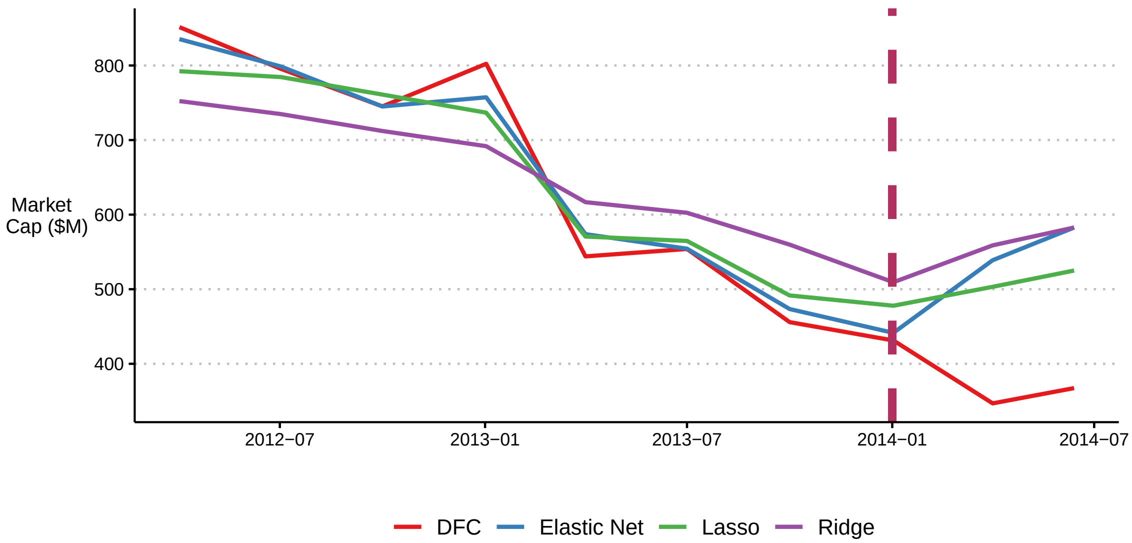

As noted in Section 4.2.2, statistical learning models tend to do better when estimated on more data. Thus, there is reason to believe that an approach that uses daily market capitalization will provide better estimates of valuation than collapsing the value to the quarterly level. In Figure 17 we show the daily percentage change in market capitalization for DFC (blue line) and its peers over the same period ending on December 31, 2013. This is the daily measured analog to Figure 15, and, as expected, shows a similar declining trend in DFC's valuation in comparison to most of its peers.

Figure 17: **Daily Market Capitalization Over Time for DFC and Peers** This figure reports the percentage change in market capitalization for DFC and the identified peer firms for the two year period prior to calendar year end 2013, through the valuation date. The data is measured daily.

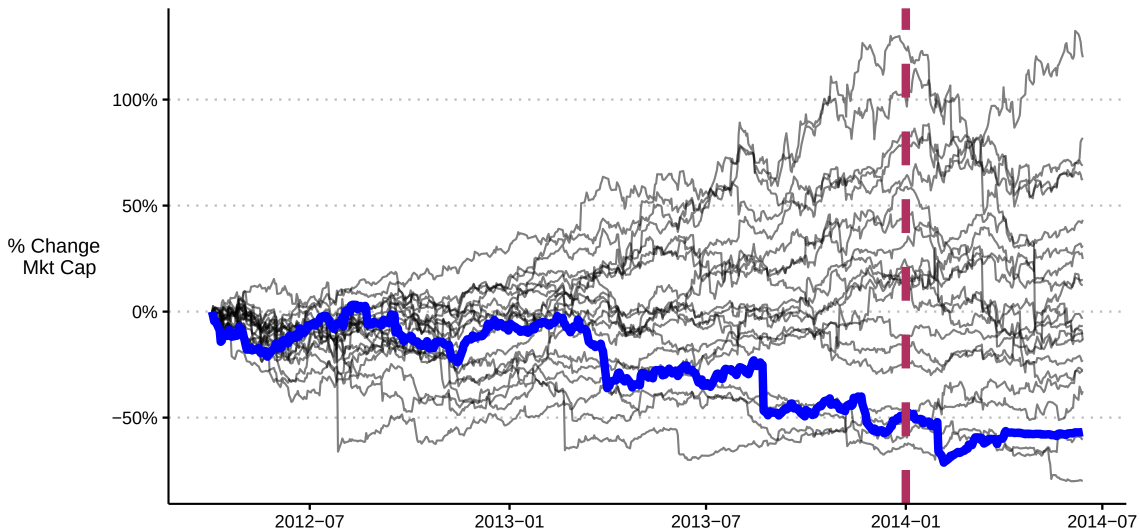

Table 9 reports the coefficient values from the penalized regression models using daily rather than quarterly market capitalization values. The outcome variable is the daily market capitalization value for DFC for the one-year period ending on December 31, 2013, and the features are the corresponding returns on the peer firms that have a complete valuation series over the relevant period. We optimize the tuning parameters using 25-fold cross-validation. The lasso model is much less sparse under this approach, with only one peer firm being dropped entirely from the model.

Table 9: Penalized Regression Weights on Peer Firms

Company	Lasso	Ridge	Elastic Net
INTERCEPT	-328.0760	-38.3504	-289.2748
ALBEMARLE CORP	-0.0403	-0.0115	-0.0497
BLOCK H & R INC	-0.0002	-0.0057	-0.0018
CASH AMERICA INTL INC	0.1240	0.1053	0.1508
CREDIT ACCEPTANCE CORP	0.1052	0.1020	0.1127
EURONET WORLDWIDE INC	0.0435	-0.0138	0.0734
EZCORP INC -CL A	0.5663	0.2943	0.5568
FIRST CASH FINANCIAL SVCS	-0.1851	-0.1071	-0.2020
GLOBAL CASH ACCESS HOLDINGS	0.5388	0.2803	0.5530
GLOBAL PAYMENTS INC	0.0511	0.0363	0.0539
GREEN DOT CORP	-0.2145	-0.1337	-0.2356
HEARTLAND PAYMENT SYSTEMS	0.0000	0.0056	-0.0389
HIGHER ONE HOLDINGS INC	0.7051	0.4946	0.7239
MASTERCARD INC	0.0032	-0.0005	0.0033
MONEYGRAM INTERNATIONAL INC	0.0425	0.0431	0.0473
QC HOLDINGS INC	-0.0030	1.4392	-0.3903
REGIONAL MANAGEMENT CORP	-0.1942	-0.0697	-0.3068
VERIFONE SYSTEMS INC	-0.0297	0.0021	-0.0317
VISA INC	-0.0066	-0.0036	-0.0066
WESTERN UNION CO	0.0131	0.0143	0.0181
WORLD ACCEPTANCE CORP/DE	0.1125	-0.0251	0.1116

This table reports the coefficient values on the peer firms using different forms of penalized regression. The outcome variable is the market capitalization for the target firm, and the features that enter the regression are the market capitalization levels for the peer firms. We use daily data for the one-year period ending in 2013, and optimize the tuning parameter using 25-fold cross-validation.

Figure 18 shows how the model predictions compare to the realized value for DFC over the estimation and valuation periods. The penalized regression models largely track the actual price series over the estimation period. The predictions exhibit less daily variation than the underlying DFC valuation series, as expected. Moreover, the predictions for the valuation date are substantially closer to the actual market capitalization after announcement, with an implied price per share of \$6.98 to \$9.74 depending on the model.

Figure 18: **Daily DFC Market Capitalization and Model Predictions** This figure reports the daily market capitalization for DFC, as well as its predicted value using different penalized regression models. The vertical dashed line represents the end of calendar year 2013.

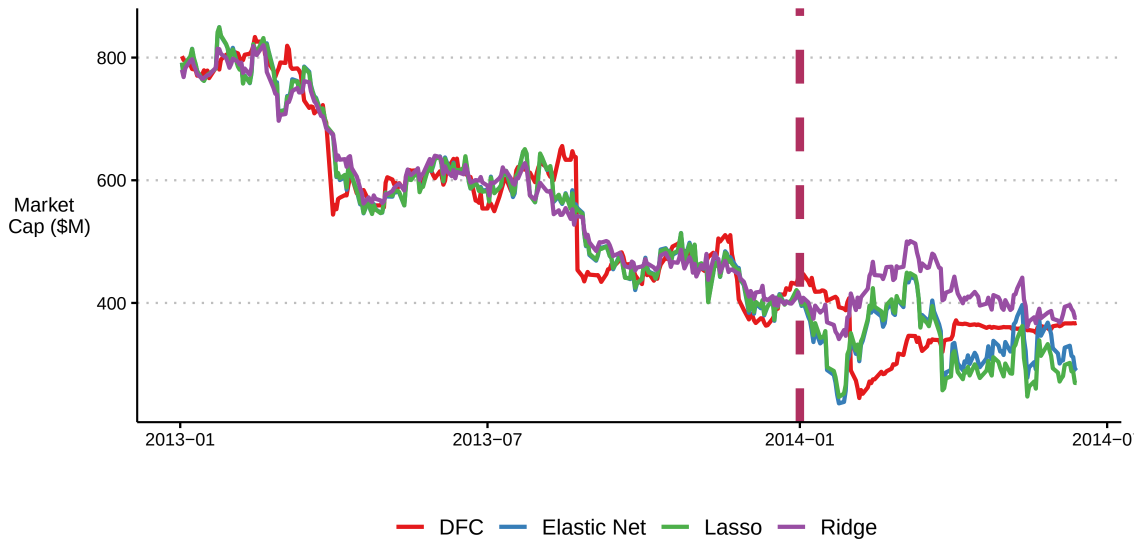

Finally, instead of using the valuation level as the outcome in our penalized models, we can instead using the returns on the stock, consistent with most modeling practice in academic finance. As shown in Section 4.2.2, these models perform better, and produce more consistent estimates of firm value over our Monte Carlo sample. The coefficient values from returns-based penalized regressions are reported in Table 10. We again see a return to sparsity, with much of the predictive power of DFC's return being driven by the market return (consistent with the Capital Asset Pricing Model), and to a lesser extent from a set of peer firms including Cash America International, H&R Block, Higher One Holdings, and World Acceptance Corp.

Table 10: Return Coefficients

Company	Lasso	Ridge	Elastic Net
INTERCEPT	-0.0022	-0.0025	-0.0022
MKTRF	0.4595	0.2053	0.4551
SMB	0.1199	0.3035	0.1571
HML	0.0000	0.2868	0.0000
UMD	0.0000	0.0363	0.0000
ALBEMARLE CORP	0.0000	0.0620	0.0000
BLOCK H & R INC	0.0655	0.0893	0.0757
CASH AMERICA INTL INC	0.3231	0.1745	0.3285
CREDIT ACCEPTANCE CORP	0.0000	-0.0020	0.0000
EURONET WORLDWIDE INC	0.0000	0.0283	0.0000
EZCORP INC -CL A	0.0000	0.0440	0.0000
FIRST CASH FINANCIAL SVCS	0.0000	0.0118	0.0000
GLOBAL CASH ACCESS HOLDINGS	0.0000	-0.0145	0.0000
GLOBAL PAYMENTS INC	0.0000	0.0202	0.0000
GREEN DOT CORP	0.0000	-0.0053	0.0000
HEARTLAND PAYMENT SYSTEMS	0.0000	0.0398	0.0000
HIGHER ONE HOLDINGS INC	0.0408	0.0591	0.0487
MASTERCARD INC	0.0000	0.0223	0.0000
MONEYGRAM INTERNATIONAL INC	0.0000	0.0311	0.0000
QC HOLDINGS INC	0.0000	-0.0280	0.0000
REGIONAL MANAGEMENT CORP	0.0000	0.0297	0.0000
VERIFONE SYSTEMS INC	0.0000	-0.0198	0.0000
VISA INC	0.0000	0.0410	0.0000
WESTERN UNION CO	0.0000	0.0482	0.0000
WORLD ACCEPTANCE CORP/DE	0.0953	0.1050	0.1038

This table reports the coefficient values on the Fama-French-Carhart factors, and the peer firms, using different forms of penalized regression. The outcome variable is the return on DFC's stock, and the features that enter the regression are the Fama-French-Carhart factors and the returns for the peer firms. We use daily data for the year prior to and ending on December 31, 2013 in fitting the model. We optimize the tuning parameter using 25-fold cross validation.

In Figure 19 we report the actual valuation for DFC in the post-affected period (black line), against the model predictions using our penalized regression approach with returns as the outcome variable. In this model, the lasso and elastic net results are functionally equivalent, as the optimal weighting between the L_1 and L_2 norm (α) is very close to 1. The results are largely consistent across penalization type, suggesting a firm value of approximately \$8.04 to \$9.19 per share.

Figure 19: **DFC Daily Market Capitalization and Returns Model Predictions** This figure reports the daily market capitalization for DFC and the prediction from a series of models using returns rather than valuation as the outcome variable. The time series for DFC is represented in black while the model predictions are in colors. The vertical dashed line represents the end of calendar year 2013.

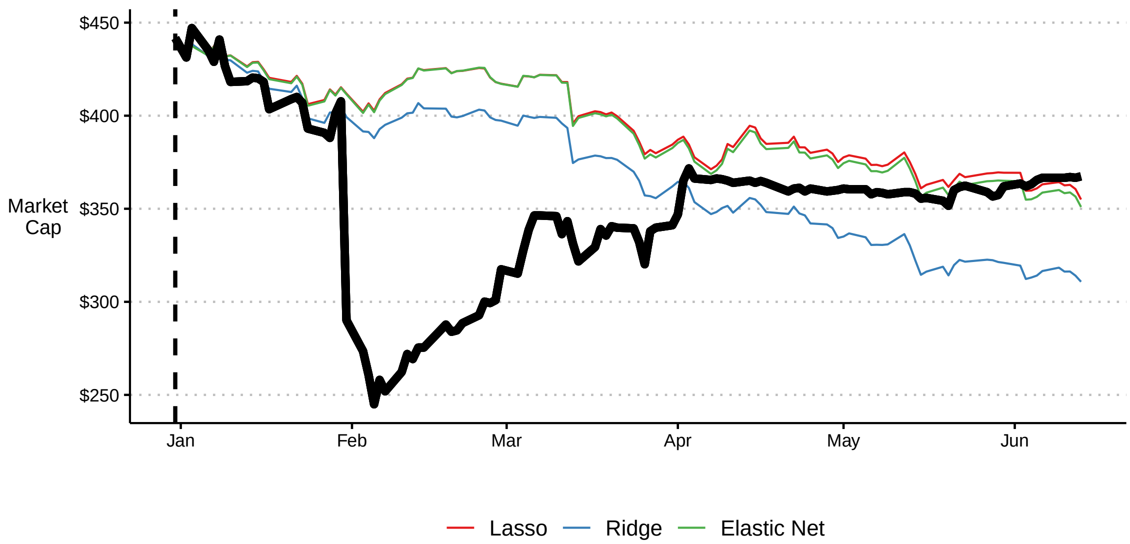

We summarize the valuation estimates from the penalized regression approach in Table 11. For each modeling choice (using either quarterly or daily market capitalization values, or daily returns) we report the estimated valuation using lasso, ridge, and elastic net regressions. In addition, we report the weighted average of the estimates, using the inverse of the cross-validation (or leave-one-out) squared error as the weighting measure. The quarterly market capitalization approach leads to the highest estimates, with a weighted average per-share value of $14.53, closer to the petitioner's than the respondent's estimates. However, as we showed in our Monte Carlo estimates, the quarterly approach produces estimates with higher variance when used to estimate value over the broad sample of publicly traded firms. Using daily market capitalization and returns as the outcome measure, we attain an estimated fair value of DFC's share price of $8.14 and $8.77 per-share, slightly higher but much closer to the respondent's estimate.

Table 11: Fair Value Estimates - DFC

Source	Type	Estimate
Petitioner's Expert		
Dages		
	DCF	$17.90
	CC	$11.38 - $26.95
Repondent's Expert		
Beaulne		
	DCF	$7.81
	CC	$8.07
Data-Driven Approach		
	Overall	$7.94
Quarterly Market Cap		
	Lasso	$13.59
	Ridge	$15.08
	Elastic Net	$15.08
	Weighted Avg	$14.53
Daily Market Cap		
	Lasso	$6.98
	Ridge	$9.74
	Elastic Net	$7.50
	Weighted Avg	$8.14
Daily Returns		
	Lasso	$9.19
	Ridge	$8.04
	Elastic Net	$9.08
	Weighted Avg	$8.77

6 Limitations and Extensions

Based on the simulation results in Section IV, we believe that our data-driven approach offers substantial benefits for the practice of financial valuation in litigation. However, as demonstrated in Figures 12 and 14, the performance of the data-driven methods improves markedly when using daily data. This works well when the targeted value is firm market capitalization, though it creates a challenge if the dispute centers around the firm's enterprise value. This is because firms only report their aggregate debt levels to the market on a quarterly basis with their quarterly 10-Q and annual 10-K filings with the SEC.

Ultimately, this should not be an insurmountable challenge for litigation valuation. Because of priority in bankruptcy, namely that equity gets paid out only after debt, the enterprise value calculation typically uses the book value of debt. As a result, for most intents and purposes, we can use the penalized regression approach to value the residual claim on equity, and carry-forward the most recent quarterly measure of the book value of debt. There is no need to value firm debt on a daily basis. To the extent that the firm issued or redeemed an anomalous quantity of debt during the contested period, one could easily enough model the firms expected debt level separately using the quarterly penalization approach derived above.

Another potential area of improvement in data-driven valuation is model-averaging. In each case we have many potential valuation estimates, depending on the choice of model and data. In other settings, model-averaging has proven quite effective at capturing the inherent uncertainty in prediction exercises where the true underlying data-generating process is unknown (CITE TO NETFLIX CHALLENGE).

To explore whether model-averaging could be useful for litigation purposes, we test its performance against the underlying models for one of our series of estimates. We focus on the use of daily market capitalization values as the outcome variable, and the penalized regression and random forest models. For each of our 10,000 Monte Carlo simulations we calculate the model averaged estimate as:

$$m = \frac{1}{4} \sum_{k} (w_k \times V_k)$$

where k denotes one of the four underlying prediction models (lasso, ridge, elastic net, and random forest), w_k is the inverse of the estimation period cross-validation squared error, and V_k is the valuation estimate for model k on date d^*. In Figure 20 we report the density of the underlying model and the model average for our Monte Carlo sample. While the model average does well in comparison to the underlying models, it doesn't appear to provide any additional predictive improvement over the best performing models.

Figure 20: **Model Averaging** This figure reports the kernel density estimates for the machine-learning approaches to valuation using daily market capitalization as the outcome variable, as well as a model-averaged estimate. The reported value is measured by the percentage deviation from the true realized market capitalization on the trading date d^* periods after the end of the fiscal quarter. The model average is the weighted average of the underlying machine learning estimates, with the inverse of the cross-validation squared error as the weight.

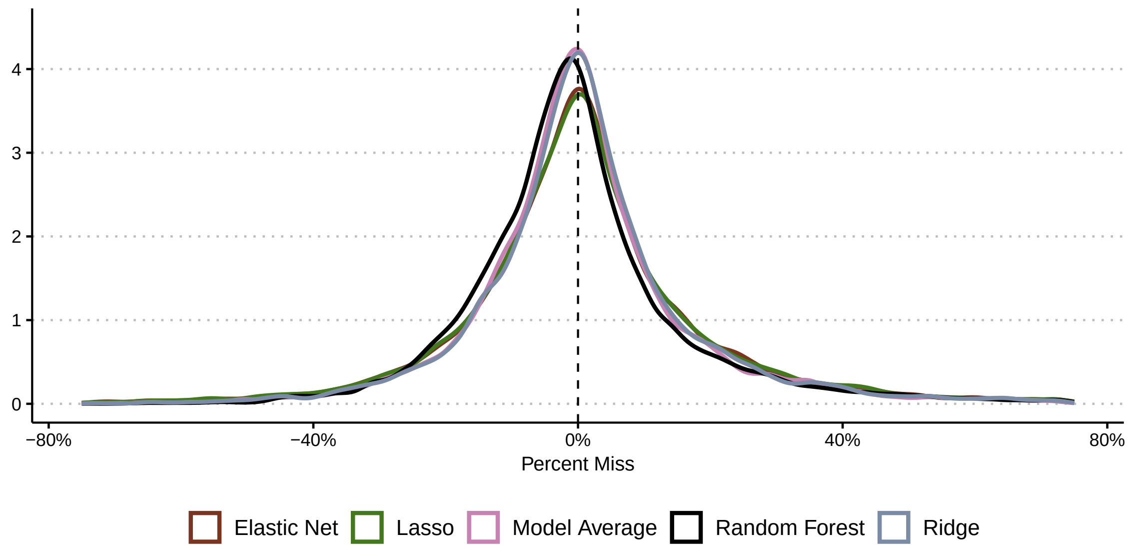

Finally, we note that our simulation approach partly hamstrings the performance of the models. In order to conduct an apples-to-apples comparison between approaches, we use the same estimation-period end date for each approach (the end of the randomly selected fiscal quarter). In any specific case, it is highly likely that there are additional, uncontaminated trading dates for estimation purposes available to the analyst. As a result, the use of daily market capitalization as the target variable should allow for more precise estimates by simply bringing the training data closer to the valutation date. In Figure 21 we report the size of the performance gains from using daily rather than quarterly data (with market capitalization as the outcome variable) by the length of time between quarter end and valuation (d^*) for each model. Unsurprisingly, the closer that the valuation date is to the end of the quarter, the more improvement is gained from using daily data.

Figure 21: **Model Improvement by d^*** This figure reports the absolute difference in the percentage miss between the daily and quarterly measured market capitalization for the penalized regression models (lasso, ridge, and elastic net) by the randomly-selected number of days post quarter-end (d^*) used as the target date. The reported curve is a smoothed generalized additive model of the difference in performance on the number of days post-quarter end, with the estimated relationship reported with a blue line, and the confidence interval for the estimate reported in grey.

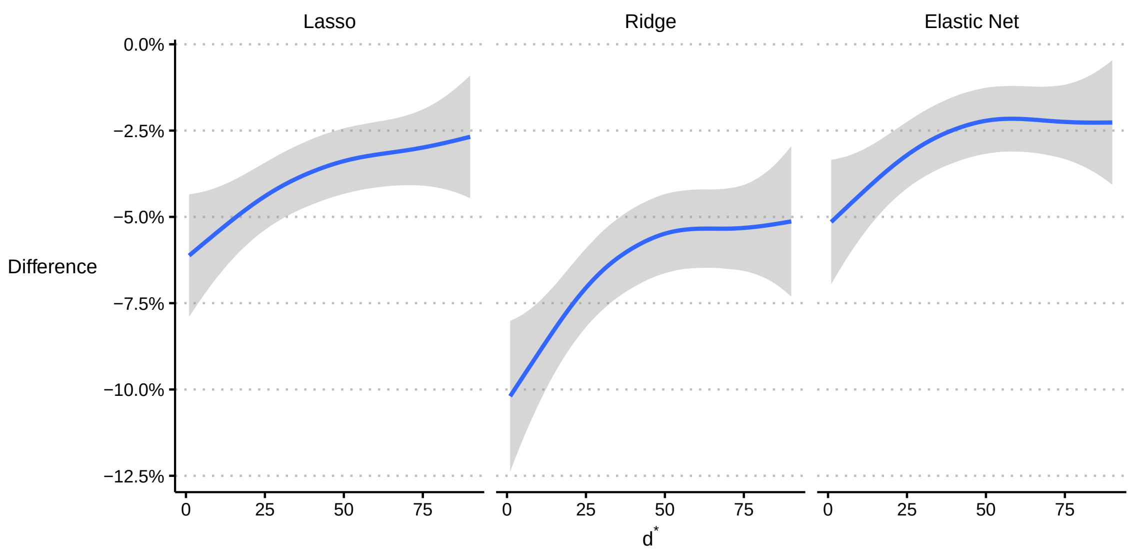

7 Conclusion

In this paper we critically examine the practice of financial valuation in litigation, highlighting the disconnect between best-practice in predictive modeling and the conventional methods used by experts. We argue that the standard approach used in for valuation in litigation creates large margins of discretion that can be used by experts to tailor estimates in a systematic manner. We propose an alternative, data-driven approach to valuation, and demonstrate through simulation analysis that these methods produce more accurate and consistent estimates. We argue for the adoption of such approaches in the legal context, emphasizing the potential to reduce the subjective element in valuation disputes and improve the alignment with modern financial theory and practice.