Peer-Review

Offutt, James, and Yale Xie. 2025. "Quantifying Legal Risk with Large Language Models: A Text-Based Investment Signal." *Journal of High School Science* 9 (3): 486–515. https://doi.org/10.64336/001c.144303

The biggest challenge I have with this paper is that - to quote financial jargon - you have bundled legal risk with a CDO (collataralized debt obligation) to create a MBS (Mortgage based security). You have then (maybe too conveniently) attributed the growth that you cannot explain (from this MBS) to your legal risk part. I am not convinced.

Let me explain in plain language:

I want to see you create an investment portfolio based ONLY on the LLM of the legal risk analysis. Then provide a comparison of the performance of this portfolio with the S&P500 (or whatever benchmark you choose). I realize that since the legal risk is not a major part of what drives stock price, your portfolio based only on legal risk may decrease over time (perhaps even go due south).

To address this - and to provide evidence that you can assign legal risk to stock performance (however small of an effect) - I would like you to create TWO portfolios using ONLY the LLM of the legal risk analysis. Portfolio 1, you would go long toward -1 and short toward +1 and in Portfolio 2, you would go long toward +1 and short toward -1. If your thesis is contributory toward investment yield, the performance of portfolio 1 should be better then that of portfolio 2 (although I would anticipate not by much)

In short, I want to see evidence (not necessarily statistically conformatory) that your LLM strategy of quantifying legal risk across concurrent quarters actually makes a difference (however small) in investment performance. If you can think of other ways to achieve this objective, please apply those as well.

_____

James Offutt

Revisions made to "Quantifying Legal Risk with Large Language Models: A Text-Based Investment Signal"

Thank you for the thoughtful review and for the opportunity to revise our manuscript. We appreciate your guidance and have addressed your request directly. Our revisions isolate the contribution of the LLM legal-risk signal—separate from "portfolio-engineering" choices—via an LLM-only equal-weight mirror test that excluded the position sizing and risk parameters included in the main model. These revisions are not only explained in further detail in this document but they are also incorporated in the manuscript via two new dedicated subsections and additional reference to the performed LLM-only test and their results are included in various other spots in the paper (including the abstract). Notably, under the "**Methods and Materials**" section, a new subsection called "*LLM-Only Test*"—where we explain the approach we took to compare the stripped LLM-only portfolio against its mirror—was added. Similarly, under the "**Results and Discussion**" section, we added a "*Comparison of Equal-Weight Portfolios"* subsection that showed the results and findings of this new LLM-only analysis. In this new results section, we added Figure 3 and Table 3 (pushing the numbers of previous Figure 3 and Table 3 to 4) which show comparison metrics and returns between the two equal-weight portfolios as well as a comparison of cumulative returns to the S&P500. The new figure and table also contain two winsorized versions for each equal-weight portfolio for further evidence as the original equal-weight strategy showed a few outsized moves that may have biased data if cleaned versions were not also overlaid.

The Reviewer's request → "ONLY the LLM" portfolios and a mirror test (normal vs. reverse)

In response to the Reviewer's instruction to construct an investment portfolio only from the LLM legal-risk analysis, then provide a clear comparison, we formed two "mirrored" portfolios that were stripped of all additional model factors outside of the trades determined by the LLM-score. To show the signal's contribution, we formed two portfolios: 1) the normal equal weight portfolio (labeled "Equal-Weight Portfolio 1") which went long toward -1 and short toward +1 (the theory-consistent direction) and 2) the mirror or reverse equal-weight portfolio(labeled "Equal-Weight Portfolio 2") which went long toward +1 and short toward -1. The goal was directional evidence that the legal-risk signal itself matters and that the returns were not random. If the legal-risk signal truly held any predictive power then Equal-Weight Portfolio 1 would yield higher (likely by a large margin when compounded) returns than Equal-Weight Portfolio 2. If the methodology was unpredictive, then the difference would converge to around 0 or the inverse would happen (Equal-Weight Portfolio 2 outperforming Equal-Weight Portfolio 1), showing that taking the other side of the the legal-risk signal's trades would've outperformed and that the model is all "noise" with no substantial signal.

<u>The following are the Implementation (what changed vs. the main model):</u>

- LLM-only weights: Removed all volatility-based sizing and nonlinear transforms. Positions were equal-weighted within each side and normalized
- The backtest was kept identical: It retained the original framework: monthly formation of sub-portfolios, three-month overlapping sub-portfolios (three sub-portfolios active after the first two months), 0.10% trading costs each way (0.20% total).
- Signal inclusions: All firms with non-zero LLM legal-risk scores from Legal Proceedings (including footnotes). Scores of 0 were not traded on by design.
- Two equal-weight portfolios (mirrored by construction):
    - Equal-Weight Portfolio 1 (normal): long negative scores (legal risk decreased), short positive scores (legal risk increased).
    - Equal-Weight Portfolio 2 (mirror/reverse): same names, same magnitudes, same period, but different sides of the trade.

<u>The following are the results of this LLM-only test (Figure 3—original and winsorized—and Table 3):</u>

- Cumulative performance (net of costs, 2020-02 to 2025-05):
    - Equal-Weight Portfolio 1: +33.76%
        - 2.5% winsorized: 26.26%
        - 5% winsorized: 30.60%
    - Equal-Weight Portfolio 2: -86.24%
        - 2.5% winsorized: -23.61%5% winsorized: -24.77%
- Portfolio-level metrics (Table 3):
    - Includes → Mean monthly returns, Sharpe (annualized), Win rate, Win-to-loss ratio, Annualized volatility, Maximum drawdown, Skewness, Kurtosis
- Mechanics: Because portfolio 2 flips the signs of portfolio 1's trades (it takes the other side) at equal magnitudes, the series are highly negatively correlated (≈ -1).
- Interpretation: The directional spread (normal minus mirror) is strongly positive throughout, indicating that decreases in LLM-scored legal risk predict relatively higher subsequent returns, while increases indicate lower expected returns—the directional and non-random evidence requested.
    - The robustness to winsorization indicates the LLM legal-risk direction drives performance beyond isolated extreme observations.

- Context relative to the main model (Table 2): Compared with the paper's flagship model, Equal-Weight Portfolio 1 shows a higher average monthly return (+1.135% vs 0.686%) but substantially higher volatility/drawdowns and a lower Sharpe. This is expected

because the LLM-only design removes position sizing and risk parameters, giving the same capital to weak and strong scores, which inflates tails (very high kurtosis) and yields more episodic outcomes. The production model's sizing and risk controls

reallocate capital toward stronger, lower-volatility names, improving Sharpe, win rate, and drawdown control while preserving the same directional signal demonstrated by the equal-weight attribution test. As expected, winsorization narrows tails (lower

drawdowns/volatility) yet preserves the normal-over-mirror ranking, reinforcing that the production model's risk controls improve investability while the underlying LLM directional signal remains the source of the effect.

- Note: The equal-weight mirror test was structurally handicapped and served attribution, not investability. The primary criterion was the directional ordering—Equal-Weight Portfolio 1 outperforming Equal-Weight Portfolio 2—rather than the absolute magnitude of risk-adjusted returns.<u>The following are the manuscript edits (where we made changes):</u>
- **Methods and Materials** section: Added "*LLM-Only Test*" describing the approach taken in this test (equal weights, normalization, the unchanged backtest framework, etc…).
- **Results and Discussion**: Added "*Comparison of Equal-Weight Portfolios*" with Figure 3 (titled "Cumulative returns of the LLM-only strategy, comparing Equal-Weight Portfolio 1 against Equal-Weight Portfolio 2, with original, 2.5% winsorized, and 5% winsorized series, 2020-02 to 2025-05, net of trading costs."). Also, we added Table 3 ("LLM-only equal-weight portfolio performance and risk metrics, including original, 2.5% winsorized, and 5% winsorized results for Portfolio 1 (normal) and Portfolio 2 (mirror)") reporting mean monthly return, Sharpe, win rate, win-to-loss, volatility, max drawdown, skew, and kurtosis for both portfolios, and interpreting these metrics in light of the attribution objective.
- **Figure/Table cross-references:** Updated the Results subsection to reference Figure 3 (cumulative curves) and Table 3 (summary metrics) together, and added a brief note explaining why equal-weight attribution is a conservative lower bound relative to the production model.
- **Abstract, Introduction, and Conclusion**: Added a one-sentence preview and mention noting the LLM-only equal-weight mirror test, that the normal equal-weight portfolio outperformed the mirror in the expected direction, and the meaning of this test outcome.

Thank you for addressing my comments. The data do seem to indicate that the legal risk metric does contribute to returns (although I am still skeptical- but that has to do with my bias, rather than to do with your data). Now that your data indicate that there is quantitative justification for your premise, I do have some more comments on the following aspects of the manuscript that I did not quite understand. Please provide more clarity in the manuscript wherever appropriate.

1.I do realize that you state "…….Firms with a legal risk change score of 0 were excluded from the trading process entirely to ensure materially informative signals….." Does this apply only to those consecutive quarters with a score of zero? If these firms did have consecutive non-zero scores at any point in time, were they then included in the long-short strategy at that period in time?

2.This long-short strategy only applies if there are consecutive non-zero numbers for legal risk in (consecutive) quarterly reports. Correct? For example, if the values were 0, 0.5, 0, -0.5, 0, 0.5, 0, -0.5…… then the strategy would be just to hold? Because there are no consecutive non-zero numbers for acceleration or deceleration of legal risk in any two consecutive quarters? Please explain in the manuscript. Similarly, if the values were to be 0, -.8, 0, -.8, 0, -.8…….the strategy would be to hold?

3.When you say “long-short”, what exactly does “long” mean in this context. For example, if the legal risk numbers in consecutive quarters were to be (say) 0, 1, 1, 1, 1, 1……..your strategy would be to hold. correct? or would it be to accumulate more stock? if yes, then would the extra amount of stock acquired depend on how the total pie is redistributed (among other stocks in the portfolio)?
4.Say I only had two stocks in the portfolio with stock A at 0, 1,1, 1,1, 1…… and stock B at 0, -1, -1, -1, -1, -1…….., and I started out with $100. Would I buy stock A with all the $100, then margin that $100 (with another $100) to borrow shares of stock B, sell the stock on the open market and then buy back stock B later, if say (next quarter), the number goes positive? (Does that not require some prediction on (legal risk) whether the number is likely to be 0, positive or negative next quarter, because the shares borrowed on margin have to be returned within a certain time?) ? Or does your model not include margin accounts? Are you dealing only with options for going short and long? Please explain in the manuscript more about how the guts of this strategy would work, with a subset of 2 or 3 stocks so that readers understand more about what is actually happening with the mechanics of the portfolio.
5.You have charted a totally new strategy with long-short based on LLM output of legal risk. Let us say that an investment bank wants to use your strategy - but they are not willing to go all the way in - they want this strategy to contribute to 5% of their current strategy (whatever that may be). How would they ‘add’ your high turnover strategy to their model? Will this not increase fees and loads?
6.Put down ‘materiality’ as a limitation. Companies are not required to disclose legal risk if that risk is not considered ‘material’, per the SEC.
7.you state “…..which were then automatically narrowed to 734 firms that contained “high-quality” SEC extractions (i.e., extractions with at least 200 characters), ensuring minimal bias and generalizability of findings……..” i would call this a deliberate selection effect; not a minimization of bias. Please discuss.
8.you state that “……..integrating real-time news feeds or litigation databases, could substantially reduce information lags,…….” Would this not bypass your LLM though; or are you implying analyzing news feeds with LLM for legal risk as well?
9.Table 1, can you add another row for the same attributes derived from a US total stock index fund (only for the loge of market value of equity) (such as the CRSP total market index fund?) This will inform the readers of how representative your sample is with respect to the total US stock index.
10.You divided 734 firms into 64 sub-portfolios= 11 firms per portfolio. Were these divided randomly? 80% of your LLM legal risk outputs are zero i.e out of 11 firms in one portfolio, you long-shorted 2 firms in one subportfolio. My question is - why did you decide to divide into sub-portfolios in the first place? Given that there are 64 subportfolios, do all your investment values represent a mean for those 64 subportfolios? if yes, then please present means and standard deviations for all investment attributes as an intermediate table. Subsequent tables will (have?) used these means for all the 64 subportfolios? Or does your paragraph beginning with “……To calculate the aggregate portfolio’s returns for July 2020, yo……..” address this question? Correct me if I have misunderstood. If I have, please provide a clearer explanation in the manuscript.
11.Do you have an explanation for why the loading on RMW may be negative (Table 5)? This does not make intuitive sense.

James Offutt

Second Revisions made to “Quantifying Legal Risk with Large Language Models: A Text-Based Investment Signal”

In response to the Reviewer’s questions and concerns, after some clarifications, below are the edits I made (correspondingly numbered to your original comments). As it came up repeatedly in the

revisions (mainly comments 1-4) , I first wanted to clarify how the LLM scoring and sub-portfolio systems work as I realized that some of the methodology was lost in translation into the writing:

How the scores are produced → For each firm and month, we ask the LLM to read the firm's current 10-Q/10-K together with its prior filing and return one number that reflects the change in legal risk between the two. We do not assign a score to each filing separately, and we do not hand-compute the change. The model compares the two filings in a single prompt and outputs a single "legal-risk change" score; that score (including zeros) is what drives all trading decisions.

How the sub-portfolios work → Sub-portfolios are time cohorts, not an equal split of firms. We run a 64-month window; each month we create one new sub-portfolio containing all firms that filed in that month with a non-zero change score. Positions opened in a given month are held for three months, so at any calendar month there are up to three overlapping sub-portfolios active.
The same firm can appear in multiple sub-portfolios over time, and—as new filings arrive—its position can flip from long to short (or vice versa) depending on the new score. Within each sub-portfolio we compute returns by averaging the long names and subtracting the average of the short names, using either equal-weights (LLM-only attribution test) or volatility-scaled weights (main model). The overall portfolio return in a month is the average of the active sub-portfolio returns.

- The explanation provided above should clear up some of the confusion involved with this question, but "Score = 0" applies only to that month's filing pair (the current 10-Q/10-K compared with the immediately prior filing). We do not assign (or store) per-quarter"levels" of legal risk and then compute differences. Instead, the LLM reads the two back-to-back filings together and returns a single change score for that month. We trade only when this change score is not 0 because a score of 0 would indicate no positive or negative change in legal severity. If the same firm later produces a non-zero ΔS in another month (so another sub-portfolio), it is included in that later month's positions. Intuitively, because we act on changes rather than stored legal risk levels, although not how the methodology technically functions, hypothetical score paths like "0.5 → 0.0" and "0.0 → -0.5" are equivalent in that both imply the score = -0.5 for that month as the

legal risk is decreased by 0.5. A zero simply means "no material change detected" in that month's comparison, so the firm is skipped for that month (consecutive non-zeros are not required for inclusion). This is already included in the manuscript.

- Same as the numbered revisions around this one, I believe there was a misunderstanding. No consecutive non-zero readings are required as this is not how the method works/ was explained in the manuscript. Each month's position is based only on that month's change score (which is assigned based only on the change in a firm's legal severity between the two back-to-back quarters). In the first sequence (0, 0.5, 0, -0.5, …), the portfolio would skip the first month (0), then short on +0.5, skip on 0, then long on -0.5, and so on (each trade belongs to the sub-portfolio formed in that month). The second sequence (0, -0.8, 0,

-0.8, …) similarly alternates skip/long/skip/long…. We have added plain-language examples in the manuscript to make this explicit and this concept is explained in various other spots in this document as well.

- Same as questions 1 and 2, the explanation provided at the top should clarify and this seems to have been a misunderstanding. We go long all firms with negative scores (legal risk decreased) and short all firms with positive scores (the reviewer's comment was implying the opposite)

and we rebalance each month. In your example 0, 1, 1, 1, 1, 1, the stock would be skipped in the first month's portfolio, then it would be shorted in each subsequent month/portfolio with a new position each time (in separate, independent portfolios). The total capital for a given sub-portfolio is set at 100% and the usage of capital in other sub-portfolios is irrelevant. We now describe this in the manuscriptaround: "In the scenario where a firm has a positive or negative legal risk change in two consecutive quarter-over-quarter observations...".

- As I tried to clarify above, here and in question 1-3, I think there was a little bit of a misunderstanding around how the scoring works. To give the short answer before the longer explanation below, while shorting requires a margin account and creates a margin requirement (account equity required to open a short trade) the model does not use margin or leverage in trades (it actually tends to not fully invest every sub-portfolio due to position sizing caps/scaling) and we do not use or mention options. In fact, as we newly mentioned in the manuscript (small note in abstract, methods, and results), employing leverage suits this model perfectly as drawdowns and volatility are low and the only shortfall is the lower-level of returns (return amplification via leverage would fix this without taking excess risk due to the model's good risk profile). I added a two-stock example to the manuscript: "For an illustrative, two-stock sub-portfolio example…".
    - For the situation of "stock A at 0, 1,1, 1,1, 1…… and stock B at 0, -1, -1, -1, -1, -1" you asked about, it would not be traded or created in the way that I think you were referring to. As explained in our **Methods and Materials** section, each LLM legal-risk change score simply quantifies the change in legal severity of a firm between two back-to-back filings (so q1 vs. q2, q2 vs. q3, or q3 vs. q4, and so on). So, the six scores ("0, 1,1, 1,1, 1" for stock A and "0, -1, -1, -1, -1, -1" for stock B) you refer to would represent six different sub-portfolios (six months since it is one sub-portfolio per month) for the six individual quarter-to-quarter legal-risk change scores that the LLM assigned. So, there would be six sub-portfolios in this example (one formed every month) with each sub-portfolio's trades being held for 3 months. So to address your other comment about "prediction of next quarter's legal risk" this is never done or needed). Also, shares borrowed via margin accounts when shorted usually do not need to be returned by a set, certain time unless something very dramatic/sudden occurs (even then it still doesn't usually occur).
    - Sticking with your example, in terms of the allocation of the $100 of that sub-portfolio, it would completely depend on the volatility during the current"regime"/time and the volatility of stocks A and B (as the model is built on the assumption that you would obviously want to put less capital towards a riskier asset). Regardless, the model never surpasses the 100% capital allowance of a sub-portfolio (so no leverage or margin for additional buying power).

    Furthermore, as mentioned in our methodology explanation, since scores of 0 are not traded on, the first of the six sub-portfolios would not long or short either stock A or B (due to the 0 scores, indicating no change in legal status). For the subsequent 5 sub-portfolios, stock A would be shorted because of its positive LLM score which shows increase in legal risk (already explained in the manuscript) and stock B would be longed because of its negative score (showing decrease in legal risk).

- The insights from legal-risk allow for additional adoption paths that reduce trading intensity while preserving the strategy's informational content (which is explained below). Furthermore, a formal explanation of this type of usage of our model is detailed in the *Strategic and Financial Applications* subsection.

    - First, managers may deploy the signal as a "watered-down" version sized to ~5% of their NAV or of the risk budget. To lower turnover, they could trade only the strongest signals—for example, names with $|LLM\ risk\ change\ score| \geq$ their score criteria (or top/bottom N by absolute score in a quintile or decile based system)—and rebalance monthly or quarterly with 3 or 6-month holds (or until the score's sign flips). Standard risk controls would apply (volatility-scaled sizing, single-name and sector caps, etc…). This selective approach materially reduces the number of trades and therefore transaction costs, while keeping the expression of the legal-risk direction.
    - Next, in terms of a research/monitoring "overlay," alternatively, the signal can function as a preliminary assessment or screener: analysts allocate due diligence on names with large negative scores (candidates for increased long conviction) and reassess names with large positive scores (candidates for hedge reduction, trims, or fundamental shorts). Existing positions can be risk-managed when legal risk worsens (tightened risk budgets, temporary hedges), and scaled when riskimproves. This mode treats the signal as input to a fundamental process rather than a standalone trading program, further containing turnover.
    - For costs and fees, computational costs are modest relative to standard research budgets. The principal concern is P&L (which the signal would likely improve) and then cost drivers like trading fees, which have fallen meaningfully in recent years and can be further minimized under the selective/overlay modes described above. Since we also incorporated symmetric 10 bps trading costs in the backtests, live implementation trading fees have already been mostly priced into the study's return capabilities.

- We have addressed your requested addition of "materiality" in the *Limitations* subsection with a paragraph starting with "A core limitation of any filing-based signal is that firms are obligated…". In this paragraph we address your revision by explaining that the signal depends on disclosures that issuers deem material under SEC rules. Thus, it may under-represent smaller or idiosyncratic legal risks and may tilt toward larger, more litigious firms. Furthermore, the subjective nature of what is deemed "material" may cause the signal to miss out on potentially influential data that did not make its way into SEC filings. We acknowledge this coverage constraint and its implications for generalizability.

- We agree that the wording was not the best as the character requirement is pretty arbitrary and could have been described better. We now describe the ≥200-character requirement as a data-quality filter that constitutes a deliberate selection: very short extracts frequently lack substantive legal content, so enforcing a minimum length improves comparability of textual inputs at the cost of excluding some filings. This was added in the area of the text starting with the following sentence: "The initial dataset was composed of 1,200 randomly selected firms…".

- Short answer: the LLM would not be bypassed and it could just use alternative data for a similar analysis, but mainly, we misworded that sentence in the previous version of the manuscript. We meant the original statement as more of a "what if" or "next steps" sentence to comment on possible future extensions of the topic (by us or other researchers). If implemented (in another study), any real-time integration we suggest would be interesting could feed those sources into the same LLM pipeline, using either rolling windows or event-to-event comparisons to generate change scores at higher frequency. Basically, the LLM would remain at the core of the method: faster inputs would reduce information lags without

bypassing the model. We fixed our misworded statement in the **Conclusion** section with the portion starting with: “Creating a similar-style model that can receive real-time news feeds or litigation database information…”.

- We agree on the value of benchmarking the cross-sectional size distribution. Because Table 1 is already space-constrained, we did not add new rows/columns to the table. Instead, we computed the same summary statistics for the CRSP Total Market Index and the S&P 500 over our sample window and integrated the comparison into the Table 1 analysis. Specifically, our sample’s natural log of market value of equity (ln(MVE)) has mean 6.999 and median 7.015, which aligns closely with CRSP Total Market (mean 7.082, median 7.128) and sits well below the S&P 500 (mean 10.428, median 10.276). We also reference the CRSP quartiles (25th: 5.539, 75th: 8.556) and S&P quartiles (25th: 9.678, 75th: 11.007) to show placement of our sample relative to each index. This addition communicates representativeness without expanding the table’s layout and demonstrated our sample’s similarity to the broader U.S. stock market but lower concentration of mega-cap and large-cap firms compared to the S&P (which was expected).

- Like the case with comments #1-3 This seems to be a misinterpretation.We clarified this revision in the manuscript and made some of the explanations more clear, but we mainly wanted to clear up a misunderstanding from the reviewer’s comment. The 64 sub-portfolios are monthly time sub-portfolio with one sub-portfolio per month (not random equal splits). Each month we formed one new sub-portfolio of all firms that received a non-zero score from the LLM that month. Positions were held for three months, so up to three cohorts overlap at any time. We (and Cohen–Malloy–Nguyen inLazy Prices) form monthly sub-portfolios and hold them for three months to align trades with filing-date information (SEC-specific), create overlapping sub-portfolios, and produce a continuous event-time return series for clean attribution and realistic rebalancing. The same firm can appear in multiple cohorts over time as new filings arrive so there is truly no spreading of the firms across 64 different portfolios. To your question about mean returns: technically, yes, return statistics are based off of the average across all of the sub-portfolios but this is only because that gives the strategy’s average returns over the entire trading duration/ sample (all 5+ years of the sample). So, it is not as if we are splitting the firms across a few sub-portfolios and then averaging those, we are pretty much just computing the average returns throughout the full duration. However, this was already acknowledged when we discussed “Mean Portfolio Return (monthly %)” in Table 2/3 and in other locations. Also, since annualized volatility is just the annualized standard deviation of the returns from the active sub-portfolios, we already report standard deviation given as volatility (standard deviation is the proxy for volatility in quantitative finance). Unless otherwise noted, reported performance statistics are computed from the aggregate portfolio time series obtained by weighting and aggregating the active sub-portfolios’ returns each month (per the procedure described in Methods). One direct spot in the manuscript we added an explanation was “(each sub-portfolio represents a specific month of trading, not a division of firms)”.
    - In the manuscript, aside from the bigger clarifications we made, we added the following two sort clarifiers:
        - “each sub-portfolio represents a specific month of trading, not a division of firms”
        - “This means that the strategy had a total of 734 unique tickers to trade in any given month.”

- We have added discussion in the manuscript noting that the strategy can load negatively starting with the sentence “The negative loading on the RMW factor indicates that the portfolio disproportionately captured returns from firms with weak operating profitability.” We explain that RMW’s negative loading indicates that the signal performs inversely with the excess return of firms with robust operating profitability over firms with weak operating profitability. In other words, our signal performs well when weak firms do better than robust firms. In the other sentence of our in-text added discussion around this we state that: “A potential reason for this is that less profitable firms are typically associated with more volatile legal controversies, so whenever a particular change occurs (whether positive or negative), they will tend to be represented in a sub-portfolio.”

Thank you for addressing my comments. Accepted.