

The dangers of taking academic evidence at face value

The Financial Times (FT) Moral Money newsletter – an influential and informative newsletter on the important topic of how businesses should serve wider society – on Wednesday 24th June 2020 highlighted recently published research¹ on remuneration consultants. Titled ‘Familiarity breeds back-scratching in executive compensation’, the piece was unequivocal about what the research found:

“What role have compensation consultants played in top executives’ ever-rising rewards? Governance watchdogs have long suspected that the people who get paid to tell companies what the boss should get paid have a vested interest in boardroom wage inflation. Now a study in the *Journal of Management Accounting Research* has backed up that suspicion with data...

... in the only language those clients understand, for every \$1,770 increase in the consultant’s fee, the chief executive can expect to take home another \$4,474.”

Why are we writing about a short piece of a few paragraphs on the arcane topic of remuneration consultants?

Is it because we think that journalists shouldn’t be writing about academic research? Absolutely not – our Centre for Corporate Governance at London Business School exists to use rigorous research to influence the practice of corporate governance.

Is it because we believe that remuneration consultants should be held above suspicion? (Disclosure: one of us is a former remuneration consultant.) Again, absolutely not. The role of remuneration consultants in the pay setting process, and the conflicts that could arise, is an entirely legitimate topic for research.

The reason we’ve chosen to write about this article is that it’s a particularly clear example of ‘academic research’ providing a veneer of respectability to a headline claim that the underlying research paper in no way supports.

This matters. We are both strong supporters of the practice of responsible business. However, to gain publicity, researchers may seek to present their work as supporting a popular narrative. Or, commentators may interpret a piece of research in a way to suit their own ends. This pattern of overstating the evidence of the benefits of a particular approach (e.g. ESG) does long-term damage. Not only does it undermine the credibility of proponents of a better way of doing business. It also leads to misdiagnosis of the problems and distracts focus away from the issues that do matter to the ones that really don’t.

By unpicking this example, we hope to help journalists and practitioners see the perils of casually recruiting research papers to their cause.

¹ *Compensation Consultant Fees and CEO Pay*, Jeh-Hyun Cho, Iny Hwang, Jeong-Hoon Hyun, and Jae Yong Shin, *Journal of Management Accounting Research*, Vol. 32, No. 1 Spring 2020 pp. 51–78

Not all journals are created equal

There is a huge range in the quality of academic studies, even those published in peer-reviewed journals. Journals vary substantially in the stringency of the peer review process, with the top tier of journals rejecting around 95% of papers. Thus, when assessing the reliability of academic evidence it is critical to pay attention to the quality of the journal in which it is published.

The Journal of Management Accounting Research is a low-ranked journal:

- The journal does not appear in the [FT's own Top 50](#) for the purpose of ranking business schools
- Even on the [Chartered Association of Business School Rankings](#) (which includes all journals) the journal ranks in Tier 2. This is the second bottom tier. Good business schools would generally disregard publications below Tier 4; the highest-quality ones would disregard publications below Tier 4*. See an explanation of the rankings [here](#).

Papers are typically published in a low-ranking journal because editors or referees felt that the paper had theoretical, methodological, or data weaknesses leading to flawed or weak conclusions. However, the peer review process is not perfect; it could be that editors or referees at higher ranking journals failed to appreciate the quality of the paper. Given the importance of the question – whether compensation consultant conflicts lead to inflated pay – we are not automatically dismissing the paper based on the low journal quality. But it is a warning sign, so we have reviewed it carefully ourselves. Our assessment is that the paper's flaws are the most likely explanation of it being published in a lower ranked journal.

Data limitations

Since 2009 listed US companies have been required to disclose, in certain circumstances, fees paid for executive compensation advice and other fees paid to the same firm of consultants where those fees exceed \$120,000. The paper uses this data over the period 2009 to 2014 to analyse the relationship between CEO pay levels and compensation consultant fees.

The circumstances driving disclosure are important, as they have an impact on the nature of the dataset and therefore the conclusions. In summary, disclosure arises in two cases:

- Where the company uses a single board-appointed compensation adviser (with no separate management adviser), in which case both executive compensation and other fees (where in excess of \$120,000) paid to the adviser must be disclosed.
- Where the company chooses to make a voluntary disclosure.

The inevitable consequence of the rules is that the only compulsory fee disclosures apply in the cases of multi-service firms. Specialist executive compensation advisers are only included in the voluntary disclosures. Once firms without compensation consultants are excluded, and observations where there are data gaps, fee disclosures arise in only 15% of firm-year observations. A little under half of these are due to voluntary disclosures. There are obvious problems with using data for voluntarily disclosing firms as they are not a random sample and there may well be other factors driving the decision to disclose. Since these omitted variables may also be correlated with the outcomes studied (e.g. CEO pay), this undermines any attempt

to draw conclusions from statistical analysis. Therefore, although results are presented for both the compulsory disclosure sample and the total sample, only the results for the compulsory disclosure sample should be considered (fewer than 10% of firm-year observations). The authors appear to acknowledge this as the headline statistic quoted in the FT paper is based only on the compulsory disclosure sample.

Given that the only reliable results are based on fewer than 10% of firm-year observations and based entirely on situations involving multi-service firms, the general applicability of any findings is questionable, as the authors themselves acknowledge:

“...sample selection issues inherent in our sample due to the fee disclosure rule make it difficult to generalize our findings to all public firms. For example, because our results on the relation between EC [Executive Compensation] fees and CEO pay are based on mandatory and voluntary disclosures of EC fees, they omit EC fee information for many firms that hire compensation consultants solely for EC services.”

In other words, the research findings do not have market-wide applicability, beyond the small subset of the data analysed.

Does the evidence show what the authors claim?

There are three major problems with the paper’s claims:

1. Identification. The authors do not identify a causal relationship from conflicts of interest to CEO pay inflation.
2. Design of Tests. Even if compensation consultants were conflicted, such conflicts would not lead to the results that the authors claim to find. In other words, the tests are not well-designed to analyse their hypothesis.
3. Significance of Results. The results are far too weak to justify the claims made.

1. Identification

The paper’s headline analysis (reported in Table 3) is the following regression:

$$CEO Pay = \alpha * Consultant Fees + \beta * Control Variables$$

(1)

They find a statistically significant coefficient on *Consultant Fees* (α). They interpret this as providing evidence for **conflicts** of interest **causing** consultants to **inflate** pay in order to win repeat business, which they term the “repeat business hypothesis”. This interpretation was picked up by the FT’s headline writers.

This conclusion is unwarranted based on the analysis carried out in the paper. In particular, the three key items in bold are not at all supported by the data, for three reasons:

- a. *Consultant Fees* are a poor measure of conflicts of interest
- b. High *CEO Pay* does not mean *Inflated Pay* (i.e. pay in excess of that which would maximise value)
- c. Even if *CEO Pay* were a perfect measure of inflated pay, and *Consultant Fees* were a perfect measure of conflicts of interest, correlation does not imply causation.

We tackle these reasons in turn.

a. *Consultant Fees are a poor measure of conflicts of interest*

The authors assume that *Consultant Fees* are a measure of conflicts of interest, with higher fees indicating higher conflict. They also assume that *Consultant Fees* are exogenous, i.e. “randomly happen” to be high or low in a particular firm. Specifically, in some firms, the *Consultant Fee* randomly happens to be high, and in those firms, the consultant has incentives to inflate pay to maintain the relationship. In other firms, the *Consultant Fee* randomly happens to be low, and there are weaker inflation incentives in such firms.

However, contrary to the authors’ assumption, *Consultant Fees* don’t randomly happen to be high or low in some firms. They are high or low for a reason (more technically, they are *endogenous*). It may be that, in some firms, *Consultant Fees* are high because the design of the contract is very complicated. Perhaps the contract needs to incentivise the CEO not just to generally improve firm value (as in all firms), but to deliver on a specific set of strategic objectives, with both financial and non-financial dimensions. Perhaps the strategic context of the firm requires an unusual type of contract, or complex board dynamics require more work from the consultant to manage the various stakeholders involved. A broadly dispersed shareholder base may require more extensive support with investor consultation. For all manner of reasons, the consultant needs to work much harder for that firm and deploy more staff.

In Table 5, the researchers try to define what the *Consultant Fee* ‘should be’ based on a regression model of firm characteristics. The extent to which the fee is above or below this amount is the ‘residual’ or ‘excess fee’. If the excess fee is positive, supposedly an incentive is created to retain that client. However, none of the firm characteristics can come close to measuring strategic context or situational complexity. These are inherently unobservable variables, while controls can only capture observable variables. This explains why their model of consultant fees has relatively low explanatory power: just 40% of the variation in consultant fees is explained by their model. As a result, even excess *Consultant Fees* could be capturing many factors other than conflicts of interest.

b. *High CEO Pay does not mean inflated pay*

The authors assume that high CEO pay is bad for the firm, i.e. conflicts cause consultants to offer contracts that are more expensive than optimal for firm value, because they are more concerned with being rehired than maximising firm value.

However, high pay can be good for the firm if it allows it to attract a high-quality CEO, or increases the chance that the firm retains a high-quality CEO. In fact, the authors themselves acknowledge on p53 the possibility that “the economic loss resulting from the client’s departure is substantial”. However, they quickly forget this possibility, and the remainder of the paper assumes that high pay equals inflated pay.

While the regression in equation (1) controls for various variables that attempt to capture what “fair” CEO pay should be, none of them can come close to measuring CEO ability. This is an inherently unobservable variable, while controls can only capture observable variables.

c. *Correlation does not imply causation*

It is so well known that correlation does not imply causation that this statement is almost a truism. Yet, no attempt is made to address the issue anywhere in the paper, nor is it even acknowledged. This would be unacceptable in any journal of reasonable quality. In such a journal, the review process would require the authors either to address the endogeneity of *Consultant Fees* (endogeneity is not a word mentioned anywhere in the paper) or, at a minimum, caveat that their conclusions may not be causal.

While the concern that correlation does not imply causation applies to most social science research, it is a particular problem for this paper, since the variables do not cleanly capture what they are meant to capture. Take the concern that *Consultant Fees* may capture the complexity of the assignment. It is entirely plausible that for companies where getting the right CEO is particularly critical, both the CEO and consultant are paid more to ensure that the contract is appropriately structured. While the authors have a long list of control variables, none of them can come close to capturing how critical it is to get the right CEO – this is an inherently unobservable variable, while controls can only capture observable variables.

Or take the concern that *CEO Pay* is related to CEO ability. For a high-quality CEO, the board may be particularly diligent about structuring her contract correctly to ensure that her incentives are correct, and so *Consultant Fees* may be high.

2. Design of Tests

Even if compensation consultants were conflicted, and there were no problems with identification, such conflicts would not lead to the results that the authors claim to find. In other words, the tests are not well-designed to analyse their hypothesis.

There are several problems with the test design. One is that it is the board – not the CEO – that hires the consultant for compensation consulting services, while it is management that hires it for other services. It is not at all clear that the board benefits from high pay; indeed, high pay may cause the board to lose a say-on-pay vote (say-on-pay was introduced in 2011 and so 2/3 of the authors’ sample is covered by it). Thus, if the consultant were primarily concerned with being reappointed rather than maximising firm value, it might recommend too low pay, to get the board past a say-on-pay vote, even if doing so leads to a worse CEO being hired, or a poorly-structured contract that fails to properly incentivise the CEO. Whether the consultant’s private incentive is to recommend lower or higher pay depends upon a complex set of factors that are very difficult to observe externally.

Thus, the tests relating CEO pay to total consultant fees are not a well-designed test of conflicts. A more relevant test relates CEO pay to cross-selling incentives (since a poorly-paid CEO may not hire the consultant for other services), but these results are weaker as we will shortly discuss. Another set of tests shows that the “effect” of consulting fees on CEO pay is stronger for poorly-governed firms. The authors hypothesise that, in poorly-governed firms, the CEO may have influence on the choice of consultant even if it is the board that formally makes the decision. However, if governance is poor, the first-order effect is that it should lead to the CEO

being better paid, regardless of which consultant is hired. However, Table 3 shows that this is not the case. Neither measure of governance quality (CEO-Chairman separation and institutional ownership) is related to pay in any of the eight specifications. This suggests that the measures are not actually capturing governance quality, potentially because they are endogenous. For example, it may be in firms that are otherwise well-governed (due to the CEO having a large equity stake in her firm), the board/investors are happy with the CEO and Chairman position not being separated. Thus, CEO-Chairman duality may not always indicate poor governance.

Another measure of governance quality, not in Table 3, is consultant tenure (Table 6). The authors find that the “effect” of consulting fees on CEO pay is stronger for firms with longer-tenured consultants. However, this does not make sense. Longer-tenured consultants are entrenched and less worried about being reappointed; it is the shorter-tenured consultants which have greater reappointment concerns. Thus, even allowing for the possibility that compensation consultants are conflicted by their reappointment concerns, the tests do not support this hypothesis. Thus, while cutting the data by consultant tenure is a valid test, the interpretation is invalid as the results actually contradict the authors’ hypothesis.

3. Significance of Results

Table 3 finds a statistically significant relationship between executive compensation fees and CEO pay. The coefficient implies that a \$1,770 increase in executive compensation fees is associated with a \$4,474 increase in CEO pay. This leads to the headline conclusion quoted in the FT article.

Even ignoring all of the above concerns on identification and plausibility of the hypothesis, these results are weak:

- As described above the models are not well-specified. This means that the chance of spurious correlations is high, and the ‘statistically significant’ results may simply be down to the quite logical endogenous factors outlined above.
- The magnitudes are small. The standard deviation of executive compensation fees is \$126,000. A fee one standard deviation above the average (around 1 in 6 fees will be this high) would be associated with CEO pay around \$320,000 higher, or around 4% of the average CEO pay of \$7,334,000. Note that a single standard deviation in CEO pay is \$5,607,000, or 76% of the average, and so even if it were a real effect, the consultant fee effect is small.
- The result is based on multi-service firms which have a mean non-executive compensation fee of \$1,651,000, compared with the mean executive compensation fee of \$177,000. Prima facie it might be expected that the cross-selling incentives for multi-service firms (i.e. recommending high CEO pay to receive higher non-executive compensation fees) would be much bigger than the repeat business incentive (i.e. recommending high CEO pay to keep the executive compensation business). Yet the regressions find no evidence of cross-selling incentives in the multi-service firms.

A pinch of salt needed

Overall, the paper provides no evidence of conflicts in compensation consulting, for the following reasons:

1. The tests do not identify a causal relationship from conflicts of interest to CEO pay inflation
 - a. *Consultant Fees* are a poor measure of conflicts of interest
 - b. High *CEO Pay* does not mean *Inflated Pay* (i.e. pay in excess of that which would maximise value)
 - c. Even if *CEO Pay* were a perfect measure of inflated pay, and *Consultant Fees* were a perfect measure of conflicts of interest, correlation does not imply causation.
2. Even allowing for the possibility that compensation consultants are conflicted, this should not show up in a relationship between consulting fees and CEO pay, nor should the relationship be stronger for long-tenured consultants. Since the CEO has more control over the consultant for non-compensation services than compensation services, the more precise test is on cross-selling incentives rather than repeat-business incentives.
3. The results are statistically and economically weak. In particular, the results are absent or in the wrong direction for the tests that are most precisely geared to the authors' hypothesis. There is no evidence that multi-service consulting firms respond to a cross-selling incentive that is around 10x higher than the repeat-business incentive in monetary terms. The results are weaker, not stronger, for short-tenured consultants who have greater reappointment concerns.

Note that these three concerns reinforce each other. No paper is perfect; even if a study is unable to precisely nail down causation, we can still learn from it if the tests are well-designed and the results are strong. However, the combination of weak results, poorly-designed tests, poorly-measured variables, and the lack of any attempt to address causation mean that the results tell us very little about potential conflicts in compensation consulting.

Journalists cannot carry all of the blame. They can't be expected to act as peer reviewers for original academic research and to critique econometric methods. Academics themselves have a responsibility not to over-claim for their research findings. But there are three questions that journalists can always ask to avoid playing a part in propagating ill-founded academic research:

- Where was it published? Was it in a top quality peer-reviewed academic journal? The [FT's Top 50](#) is a decent starting point.
- Do the researchers have an incentive to interpret a particular finding from the research in order to make it newsworthy?
- Have the researchers demonstrated causation, or is it really just correlation? Are there plausible alternative explanations for the result?

Academic research has a vital role to play in influencing business practice for the better. But only if it is evaluated and interpreted sceptically.

Professor Alex Edmans
 Dr Tom Gosling
 LBS Centre for Corporate Governance

1st July, 2020