

Learnings From 1,000 Rejections

Alex Edmans
London Business School, CEPR, and ECGI

Current draft: April 9, 2023

Abstract

The *Review of Finance* aimed to significantly increase its standards over my 6 years as Managing Editor and 1 year as Editor. To comply with these new standards, I had to reject nearly 1,000 manuscripts. This article aims to use these rejections constructively by distilling common reasons for rejection to guide future research. They are divided into three categories: contribution, execution, and exposition. Beyond extracts from decision letters that give reasons for rejection, this article also shares excerpts that shed light on the editorial process, such as how an editor weighs up feedback to reach a decision, as well as emails to authors outside formal letters in response to queries on the process.

Email: aedmans@london.edu, London Business School, Regent's Park, London NW1 4SA. I thank Jonathan Cohn and Josef Zechner for helpful comments.

1. Introduction

Over 6 years as Managing Editor of the *Review of Finance* (2017-2022) and 1 year as Editor (2016), I handled 1,058 manuscripts.¹ At the beginning of 2016, when Franklin Allen was Managing Editor, the *RF* started to apply strict top-three standards; we formally stated this in an editorial the following year (Allen and Edmans, 2017). As a result, I unfortunately had to reject 998 of those papers; the remainder were either accepted or are still going through the editorial process. I understand the sting of rejection, having experienced it many times myself, and am sorry to the author teams who chose the *RF* as a potential outlet for their paper but we were unable to publish it.

The purpose of this article is to use these rejections constructively, by learning from them to guide future research. In particular, there were a number of recurring themes across many of the rejections, and I will use anonymized extracts from my rejection letters to provide concrete illustrations of these themes. Moreover, 517 submissions were desk-rejected (with a further 41 desk-rejected after receiving a screening report, typically from an Associate Editor), principally due to not studying an important enough question. Learning what types of questions are unlikely to cross the bar will hopefully save other researchers the substantial time and effort required to write a paper when the topic may have limited promise to begin with.

These suggestions are divided into three themes: the contribution, execution, and exposition. The final section shares extracts on issues other than how to improve a paper, such as explaining how editors might weigh up different factors in reaching a decision. In addition to decision letters, it also contains excerpts from my replies to author requests about the editorial process, for example whether they could resubmit a rejected paper because they believed they could address the referee's concerns. The goal is to shed light on aspects of the editorial process that otherwise might be opaque.

¹ 1,058 is the number of unique manuscript IDs. Resubmissions of a manuscript retain the same ID and thus do not add to the count.

I stress that these extracts are only from my decision letters and emails and thus my personal opinion on what makes a publishable paper. Other journals, future editors of the *RF*, and even co-editors during my tenure at the *RF* may well have different views. In addition, I largely handled corporate finance papers (plus some behavioral asset pricing); the common reasons for rejection of non-behavioral asset pricing may be different.

I recognize the risks in sharing extracts from decision letters. First, some readers may disagree with my comments, either those that led to rejection or those that were offered constructively to authors even if not pivotal for a decision. Many comments will be at least partially subjective and people can legitimately have differences of opinion; shortly after my appointment I introduced three-year terms and a two-term limit to ensure cognitive diversity. I will have made mistakes; in one case I rejected a paper for insufficient contribution, then saw the significant impact it was having and invited the authors to send a new version, which was refereed afresh and eventually published. There is competition between journals, and one that keeps making mistakes will lose good papers to its peers. An editor's criteria can never be completely objective but I can at least provide transparency on what these criteria were. Second, like the Unabomber, I might be identified by my idiolect – it may be that certain phrases in my decision letters were also used in my referee reports for other journals. While reviewers may have borrowed some phrases from my decision letters for their own reports, or come up with these phrases independently, seeing certain language may increase the likelihood that I was the referee. If this is the case, I hope that I treated the paper fairly, regardless of the outcome.

In what follows, I will use female pronouns for editors and authors and male pronouns for referees and readers. Where there would otherwise be confusion, I will use “manuscript” to refer to the paper under review, and “paper” to refer to other papers that the manuscript cites (where there is no confusion, I will use “paper” for the former). I also endeavored to remove all specifics to preserve the

anonymity of the submitted papers, hence often referring to a paper as being about the effect of X on Y . Sometimes the extract will need to contain some details (e.g. its topic) in order to make my point, but I will only include such information if it is still sufficiently general that it will not identify the paper. Where possible, I changed the actual topic of the paper to a different one if doing so still allowed me to make the same point. Similarly, I have changed page numbers, footnote numbers, and variable names, and paraphrased some elements.

This article is related to other papers providing advice on the finance profession. Examples include Pedersen (2021) and Weisbach (2021) on the profession in general; Edmans (2022) on pursuing a purposeful career; Cochrane (2005) on academic writing; Spiegel (2012), Hirshleifer (2015), and Berk, Harvey, and Hirshleifer (2017) on how to improve the refereeing process; and Butler and Crack (2022) on the academic job market. Some of the above articles are from former editors and concern the referee process; their aim is to improve the quality of refereeing. This article provides a window on an editor's communication with authors; its aim is to improve the quality of research.

2. Contribution

2.1 Results are Insufficiently Novel

You are clear that you are the first to study the effect of X on Y . However, [citations] have already shown that X affects Z , and [citations] have already shown that Z affects Y . Thus, your paper is a convex combination of results already known in prior literature. Given that we already know that X affects Z and Z affects Y , it is not too surprising that X affects Y .

You are clear that you are the first to study the effect of X on Y in the UK. However, [citations] have already shown that X affects Y in the US. Unless there are plausible reasons (such as different institutional features) for why the US results would not automatically extend to the UK, a reader would not Bayesian update much after reading your paper.

We already know that X affects investment. You cite some papers, but there are others with more precise identification strategies, such as [citations]. While these papers focus on capital expenditure, it's not too surprising that the effects also extend to R&D.

We already know, from a large literature that you faithfully cite, that there are peer effects in many other corporate policies, so it's not surprising that there are also peer effects in X .

The four above paragraphs capture the same concern – even if a result has not been shown before, if the reader would have expected the result given prior literature, the contribution from demonstrating it is small. While there might still be a non-zero contribution from explicitly documenting the result, it is unlikely to be strong enough for a top journal.

The first extract is self-explanatory. The second describes an extension to another region, the third to another type of investment decision and the fourth to another corporate decision. The same concern applies to other types of extensions, such as from one industry to another, or one crisis to another. Similarly, if a result has already been demonstrated in general (e.g. across all industries), documenting it in a specific setting (e.g. in one industry) is unlikely to be a publishable contribution.

In my first year as an assistant professor, I discussed one of my fledgling papers with an editor in an office visit. He was unimpressed, saying that my paper was telling him his prior. He said that, if a reader spends a couple of hours going through your paper, you should change his prior – teach him something new that he didn't know, or could not have guessed, beforehand. While he didn't use these words, another way of phrasing his point is to ask – does a paper increase social welfare? If the benefit that a reader obtains from reading the paper is less than the cost of doing so, it may not. Thus, it is insufficient for a paper to make a strictly positive contribution; it must make a significant one.²

Note that changing your prior might involve strengthening your prior, not just overturning it. It's not the case that only one paper can ever be written demonstrating one particular result. It could be that an earlier result was found on a small sample, or using an imprecise identification strategy that couldn't quite rule out alternative explanations. A subsequent paper using a much larger sample or a

² I applied a constant cost of reading papers across all manuscripts, regardless of their actual complexity, and thus the bar for contribution was the same. However, if I issued an R&R, I then tried to reduce the paper's complexity and length. Another approach might be to vary the bar with the complexity and length of the paper.

more precise identification might strengthen our prior and make a significant contribution to knowledge.

2.2 Results are Insufficiently Important

You are clear that you are the first to study the effect of X on Y. However, to be published in a top general-interest journal such as the RF, it is not sufficient for a paper to be novel – a paper must also be important. Unfortunately, it is not clear why the effect of X on Y is of first-order interest. As you acknowledge on p5, many papers have studied other determinants of Y, and X seems to be “just another” determinant of Y to add to the long list of other determinants that we are already aware of. It is unlikely that a survey paper on the determinants of Y will mention the effect of X, or a corporate manager deciding Y will pay much attention to X.

The main result seems somewhat of a curiosum (echoing the referee’s “factoid” comment) – a temporary shock which may not make a major contribution to the Y literature since it’s a one-time shock. A survey paper on the determinants of Y might have a section on [one determinant], and that section might include [list of sub-determinants], but your paper may not get any mention at all, or only a passing mention, since it’s not that relevant outside [the temporary shock].

We already know that X leads to many positive outcomes, such as greater firm value, profitability, productivity, and innovation. Y seems to be “just another” outcome of X to add to the long list of outcomes that we are already aware of. That X also leads to greater [financial variable] seems second-order compared to the other outcomes, in particular since it is unclear whether such a shift is positive or negative for firm value. If it is positive, then your paper simply reinforces the results of prior literature, that X is positive for firm value. If it is negative, then it is unlikely to change our views on the desirability of X. A policymaker contemplating whether to encourage or discourage X would likely base her decisions primarily on its effect on the above real variables than on [financial variable]; your result is like the proverbial rabbit in a horse-and-rabbit stew.

But why do we care about this result? It’s true that you document a positive effect of X which, prior to your paper, I thought of as being unambiguously negative. However, the small benefit that you document is unlikely to offset or even mitigate in any non-trivial way the negative effects of X. At present, I fear your paper risks being seen as a “curiosum”, i.e. finding a consequence of X that’s “curious” in that it’s positive unlike all of the negative effects previously documented, but not important/interesting compared to all the negative effects. Indeed, on p2 you write that a contribution of the paper is “We also add a new wrinkle” – I know that you did not intend it that way, but the term “wrinkle” is rather unfortunate as it suggests that you add only a small nuance to a well-documented existing phenomenon.

I understand that views of importance are somewhat subjective, so let me try to provide the thinking behind my view. Some determinants of Y are important because they help us evaluate major theories of Y (e.g. [list of theories]), but this is not the case here. Other determinants may be of interest because they teach us about how financial markets are not just a sideshow but affect the real economy (e.g. on how equity market valuations affect Y), or because they are factors that a company can affect (e.g. finding that X_1 affects Y is interesting, because companies then know that they can increase their

Y by changing X_1). In contrast, that X_2 affects Y seems to be more of a factoid rather than something that allows us to evaluate major existing theories, or teaches managers how they should determine Y . It is out of managers' and policymakers' control so it's not clear what we would do with your result.

In the above extracts, the concern is not that the result could have been predicted from prior literature, but the result is simply not important. While importance is obviously subjective, whether a finding would likely be included in a survey paper is one potential litmus test. Some authors conduct “research by matrix”, where the rows (columns) correspond to potential X (Y) variables and they try to find empty cells. They scour the literature and find that it has shown the effect of X_1 on Y_1 , X_1 on Y_2 , and X_2 on Y_1 , but not X_2 on Y_2 . However, it may be that no-one has studied this question because it is not interesting. If the paper finds “just another” determinant of Y_2 , or “just another” effect of X_2 , to add to the long list of current determinants and effects, it may not be first order.

I'm afraid the magnitude of the contribution is not strong enough for the RF, as we already know from a large literature (that you dutifully cite) that X leads to Y . You write that you identify the mechanism for why X leads to Y , but this is more an extension of existing research rather than a major contribution in its own right. Often it is useful to understand the mechanism through which X leads to Y as this affects the interpretation of the results – for example, if a paper found that earnings increased, a subsequent paper could study whether the increase is due to greater efficiency or greater myopia as this affects whether the increase is desirable. However, all the potential channels through which X leads to Y are consistent with efficiency. Thus, showing that X leads to Y through channel A rather than channel B does not change our view on the desirability of X .

Documenting channels through which previous results operate is useful if the channels affect the interpretation of the results, but this was not the case here. Unfortunately, the manuscript delivered a “Section 6.2 result” which might be an interesting add-on at the end of a paper documenting that result, but was not a publishable contribution in its own right.

Below is an extract from an R&R letter which explains a situation in which documenting channels is interesting.

In general, at the RF, we try to avoid over-refereeing and unnecessary “nice-to-have” extensions; indeed, the report is commendably short and so I don’t need to do the paring down that I typically do with many reports recommending R&R. However, I do agree that some evidence on the mechanism is important for the paper to be publishable – otherwise, your paper documents a result but not why this result may be the case. Thus, it’s difficult for the reader to know what to take away from your paper and may see it mainly as a curiosum. I recognize that it’s unlikely that you’ll be able to precisely nail the mechanism and this is not the bar for publication, but some suggestive evidence would be important, and the referee provides constructive suggestions in this regard.

2.3 Topics Does Not Fit a General-Interest Finance Journal

The paper is well-motivated and the results make sense, but it will have greatest impact in an accounting journal. [Topic of paper] is important issue in accounting, and indeed most of the most closely-related papers that you reference are accounting papers. However, I’m afraid that [topic] is of less interest to a finance audience. Thus, if your paper were published in a finance journal, I fear it would not have the impact it deserves. I tend to err on the side of having a broad definition of finance, since I’m aware that, to move up in the rankings, a journal needs to take papers on non-traditional topics, since a good paper on a standard topic would likely be accepted in a top-three journal. However, even under this broad definition, I’m afraid that [topic] falls out of the scope.

Your paper is thoughtfully written and executed, has a clear hypothesis and plausible results. However, it would be a much better fit for a top organizational behavior or management journal, as this paper doesn’t really fit a finance journal. To be sure, papers on X have been published in top finance journals, but these are papers that link X to finance outcomes, such as profitability and firm value. A paper that studies the determinants of X is a much better fit for journals where X is the “end goal”, whereas finance journals have either real variables (e.g. productivity, innovation) or financial variables (e.g. stock returns, dividend policy, leverage, compensation contracts) as the “end goal”. By analogy, I’m sympathetic to the importance of employee satisfaction, having worked in the area myself, but a paper on the effect of hybrid working on employee satisfaction would not fit a finance journal.

Your paper is technically sophisticated, but I’m afraid that it’s too niche for a general-interest journal. It would be a better fit for a top field journal focused on [topic]. While your paper would likely be of great interest to readers also studying [topic], it is unlikely that the general-interest reader will read it. A reader seeing the title of your paper is unlikely to read even the abstract.

The paper has a clearly defined research question, but I’m afraid that X is too niche a topic to warrant publication in a top general-interest journal such as the RF. Many readers might not know what X is in the first place.

The first two concern a paper’s fit for a general-interest *finance* journal – that it needs to be of interest to a finance audience. Sometimes, papers might be squarely on accounting, operations research, organizational behavior, or macroeconomics. One sense-check before submitting the paper

is the proportion of papers in the bibliography that are in finance journals; if this is small, the paper is unlikely to be a good fit for one. Note that authors should not artificially pad the bibliography with finance papers to dupe editors and referees into thinking the paper is on finance; they will easily see that it is not a good fit from the content. Instead, the sense-check is useful for the authors – is the bibliography that they naturally came up with predominantly finance?

The second two paragraphs concern suitability for a *general-interest* finance journal. Here, there is no doubt that the paper is on a finance topic; instead, the question is on whether it of sufficiently broad interest to be publishable in a general-interest journal. While this is more subjective than whether a paper is on finance vs. accounting, the bibliography check remains useful – if most of the references are to field journals, a field journal may be a better fit. The “abstract” question may also help; my hope was that anyone picking up an issue of the *RF*, or receiving the Table of Contents email from the European Finance Association giving notification of a new issue, would be willing to read the abstract of every paper in the issue regardless of its topic. While it may only take 30 seconds to skim an abstract, this is a surprisingly difficult bar to cross.

The final paragraph is a more elementary check – if the topic of a paper is something that the average reader has never heard of, then it is unlikely to be of general interest. Note that this should not be a hard-and-fast rule since a paper can make a contribution by informing the reader about something he *should* have heard of. However, such papers are clear why the topic is relevant for a general finance audience – *why* it should hear about it. The first paper I ever accepted (as editor, before I became managing editor) was Yermack (2017) on blockchains. While many readers may not have heard of blockchains when I solicited it in 2015 (I admit I confused them with blockholders when I first saw the title), the paper explained their relevance to general corporate governance.

2.4 Results are Insufficiently Generalizable

The paper has a clear hypothesis and plausible results, but I'm afraid an analysis of a single case is not sufficient to draw conclusions about the effect of X on Y in general. The external validity is limited – it may be that Y responds differently in different cases. To be sure, Q is an important event and impactful papers are written on Q alone, but the topic of such papers is Q rather than trying to form more general conclusions (e.g. on the effect of X on Y in general).

Here, the concern is not that the topic *Y* is of sufficient interest to a general audience, but the setting *Q* in which *Y* was being studied was a single event. Note that specificity alone is not necessarily sufficient for rejection. For example, studying the effect of *X* on *Y* during event *Q* may be publishable if there are logical reasons for why the findings might be generalizable to other events – for example, if *Q* is a recession and the results would likely extend to other recessions. Alternatively, even if *Q* is unique, the paper might be publishable if *Q* is a particularly important event. Here, the concern is not that the event *Q* was unimportant, but that the paper's research question was the effect of *X* on *Y* in general. Editors try to find a path to publication – to think about whether a paper would still be publishable if it narrowed down its scope or toned down its claims. However, in this case, if the paper narrowed its research question to the effect of *X* on *Y* during event *Q*, its contribution would be too small.

2.5 Paper Considers Only One Side of the Trade-Off

It is logical that X would increase Y; it is difficult to think of a world in which this would not be the case. The big question is whether the increase in Y is worth it compared to the cost of X. Thus, you don't actually answer the question of whether X creates value – even if it increases Y, it could reduce value, if the benefits don't offset the costs. For a journal with a less high bar, documenting the benefit alone might be sufficient for publication, but unfortunately this is not enough for publication in the RF.

Nearly every financial decision involves both costs and benefits. Documenting only one side of the trade-off can be a publishable contribution, if that cost or benefit was not obvious. However, this paper's research question was on whether *X* overall creates value, and it only studied the benefits. If

it scaled down the research question to focus on the benefits alone, it would not be publishable since the benefit was not surprising.

2.6 Paper Lacks Clear Hypotheses

I am very open to behavioral factors - even seemingly “wacky” factors - affecting asset prices. However, I’m afraid the a priori hypothesis for why X should affect asset prices is weak, even for a reader like me who’s sympathetic to the idea that markets might be inefficient. While there is evidence in the psychology literature that the other behavioral variables that you cite have strong effects and thus may ultimately feed through to asset prices, this does not seem to be the case for X.

I’m afraid there was no clear theoretical reason for why X should increase the efficiency of Y. You provide many reasons for why X might increase the level of Y, but not the efficiency of Y. Indeed, X might increase Y beyond its optimal level.

You have a clear hypothesis for why X should affect Y₁, but you study Y₂ and it is not clear how Y₁ affects Y₂ – it could catalyze it, or it could crowd it out (Y₁ and Y₂ could be complements or substitutes). Thus, you do not have a clear hypothesis on how X should affect Y₂, and so it’s difficult to know what we learn from your findings.

You have an interesting dataset, but unfortunately the current version of the paper does not ask interesting questions with it. It is not clear what hypotheses you are testing with the data and thus what the reader learns from the results. Some of the tests simply correlate various variables together with no specific directional hypotheses given. Are you hypothesizing a specific direction of association, or are differences in either direction enough for you to claim victory?

You state that “The paper’s primary contribution is to [document heterogeneity]”. However, it is not clear why this heterogeneity is interesting. The [description of one type of heterogeneity] is indeed interesting because this is predicated on theory. However, most of the other tests seem to be rather kitchen-sink. You don’t form any hypotheses, but seem to simply regress Y on whatever data you happen to have available and split the results in different ways. But what do we learn from the fact that [the effect is stronger in sub-sample 1 than 2]? How does this change our view of the world? Would our view of the world be any different if you found the opposite result? Was there a specific directional hypothesis that you were testing with this cross-sectional analysis, or would differences in any direction have been enough to “claim victory”?

The first extract concerns the strength of the hypothesis. There is a potential story for why X may affect Y, but the channel is sufficiently weak that, even if there was a strong correlation in the data, it is likely to be spurious. Having a strong hypothesis, not just a hypothesis, is important since researchers can almost always reverse-engineer a hypothesis after finding a result. Given researchers’

incentives and ability – given computing power and datasets – to find significant results (Harvey, 2017), it is essential for authors to convince the reader that they formed their hypothesis before studying the data.

The second extract concerns the precision of the hypothesis. The hypothesis concerned the efficiency of Y , but all the arguments provided concern the level of Y , and could plausibly lead to Y increasing beyond its optimal level.

The third extract also concerns precision, but is more intricate. Here, the hypotheses concerned the effect on Y_1 , but the authors actually studied Y_2 . This is not a problem if there is a clear link between Y_1 and Y_2 . However, here Y_1 and Y_2 may be complements or substitutes. Often authors will argue “it’s an empirical question” as a defense to unclear hypotheses, but this defense is often weak. Certainly, if the goal of the paper is to study whether Y_1 and Y_2 are complements or substitutes – i.e. this is an important question in and of itself, because there are theories or economic hypotheses in both directions – then, empirically answering this question is interesting. However, the goal of the paper was to study whether X affects Y_1 . Finding that X increases Y_2 could indicate that X increases Y_1 and Y_1 and Y_2 are complements, or that X decreases Y_1 and Y_1 and Y_2 are substitutes.

The fourth and fifth extracts concern the direction of the hypothesis. Some papers simply regress Y on lots of different X variables without hypothesizing particular directions, or compare the relationship between X and Y across different subsamples (e.g. large firms vs. small firms) also without a clear directional hypothesis. This makes it difficult to know what to take away from the results – what they teach us about the world. Note that unclear hypotheses are different from conflicting hypotheses. The latter arises for when hypothesis A (B) predicts a positive (negative) relationship; then, finding a positive association rejects B but not A and moves our prior. The former exists when there are no hypotheses for any link in any direction.

As with many points in this article, the above principle is not a hard-and-fast rule. Sometimes, if the authors are introducing a brand new dataset, simply documenting correlations is interesting; future research might help us understand what the correlations teach us about the world. However, if a paper is working in a mature field, running regressions without a clear set of hypotheses is unlikely to be publishable.

3. Execution

I regularly rejected papers due to issues with execution. However, such issues were often specific to a particular paper and described in the referee report. There are relatively few general themes, so this section will be comparatively short.

3.1 Instrumental Variables

You mention that you use instrumental variables in the introduction, but you do not explain what your instrumental variables are. It is very difficult to find valid instruments in a corporate finance setting and I am sorry to say that some referees will view claiming “we have instruments”, without explaining what they are and why they are valid, to be trivializing the identification process – finding valid instruments is very difficult, and the burden is on the author to justify why the instrument satisfies the relevance criterion and exclusion restriction. Indeed, some readers get a bit suspicious when a paper claims an IV approach in the introduction but does not describe the instruments, as it suggests the authors may not be sufficiently convinced about the validity of the instruments themselves, and so do not wish to describe them in the introduction but bury them deep into the paper.

The above paragraph is actually on the exposition rather than the execution. However, since it was often immediately followed by a paragraph describing concerns with the instrument itself, I include it here. I now give examples of such paragraphs.

When we finally get to your instruments on p25, we learn that they are unfortunately invalid. One instrument is the industry average of X . However, using peer group averages as instruments is nearly never valid since any omitted variable at the individual firm level is simply soaked up at the group level (see Gormley and Matsa (2013), Section 2.3.4). While it is true that some papers have got away with using peer group averages in the past (as you cite on p26), it is now well-established that peer group instruments are invalid.

The last sentence highlights a common theme across many rejected papers – they often appeal to other papers to justify a certain approach. Often these papers are unpublished or in minor journals. Even in major journals, it may well be that research has evolved since then, and a methodology that was considered acceptable in the past is now considered to be flawed. Certainly, it is useful to cite papers that use a similar approach if the instrument is still considered to be valid – the earlier papers may have done the “heavy lifting” describing the institutional details behind the instrument and it is not necessary to repeat every single detail again.³ It is also important to acknowledge the paper(s) that first employed the instrument. However, just because other papers used the same instrument does not mean that it is valid.

Your second instrument is lagged X. However, valid instruments come from outside the system; this instrument is within the system and thus affected by the same omitted variable and reverse causality concerns as the contemporaneous variable. For example, the firm could have anticipated this year's Y and chosen last year's X in anticipation (reverse causality). Alternatively, it could have anticipated this year's W [an omitted variable] and chosen last year's X in anticipation (omitted variables).

You then attempt to justify your instruments with statistical tests, but there is no way to test for instrument validity – see Section 3.4 (“So-called tests of instrument validity”) in Roberts and Whited (2013). An instrument is invalid if correlated with the error term. Since the error term is unobservable, this correlation cannot be tested. More specifically, you conduct a test of overidentifying restrictions. However, such a test compares the relative validity of two instruments Z_1 and Z_2 . It only tests that Z_2 is valid conditional upon Z_1 being valid (which can never be proven). The test will thus give a pass not only if both instruments are equally valid, but also if both instruments are equally invalid.

These two extracts are on different issues and hopefully self-explanatory.

I saw a very good discussion which mentioned Sun Tzu's “The Art of War”, which argues that the first step in war is to “Know Your Enemy”. Similarly, the first step in addressing endogeneity is to understand very precisely what the endogeneity problem is. Only once you've diagnosed the problem can you explain why your suggested solutions to the problem are effective ones. Similarly, only once the reader can understand the problem can he/she evaluate whether the solutions used are effective, and thus how strong a causal inference to make from the paper.

³ On the other hand, a paper should be self-contained and the essential details should still be described in the paper.

While strictly a comment on the exposition rather than the execution, it is related. This paragraph often preceded a paragraph explaining why the instrument failed to address the specific issues plaguing the paper. Other times, I made the comment for a quite different reason – there was actually no clear endogeneity concern, but a prior referee may have raised an issue to signal competence to the editor (as modeled by Hirshleifer, 2017), or thinking that a longer report would be more helpful to the authors. This problem exists far beyond endogeneity; sometimes, authors will conduct a robustness check without explaining what the concern is, or verify robustness to an alternative measure of a particular variable when there was nothing wrong with the main measure. In such cases, I typically asked the authors to remove the analysis and shorten the paper.

3.2 Log-Transforming Count Data

In equation (11), you have $\log(1+Y)$ being the LHS variable. I know many papers do this when there are lots of zeros, but I don't think this is correct. The coefficients on the RHS variables are impossible to interpret – with $\log(Y)$, you can interpret them as a percentage change, but with $\log(1+Y)$, they have no interpretation. There does not seem a logical reason why the RHS variables should affect $\log(1+Y)$, whereas with a level variable there is at least a logical relationship (a linear one). Moreover, it is arbitrary that you choose to add 1 before taking logs: you could alternatively add 0.1 or 2 and get different results, whereas adding a constant to a level variable will not affect the estimation of the coefficient. Can you just have “Y” being the LHS variable? See Cohn, Liu, and Wardlaw (2022) for further description of the problem and proposed solutions.

This comment is another example of common practice later shown to be incorrect. No paper was rejected due to this practice as it is fixable, but I sometimes provided this comment as a suggestion. Cohn, Liu, and Wardlaw (2022) focus on the case in which the dependent variable is $\log(1+Y)$ because there is a natural solution of a Poisson regression. However, the same problem occurs when the independent variable is $\log(1+X)$ – the addition of 1 is arbitrary, and the coefficient has no interpretation, although there is no natural solution other than having X as the independent variable. Before Cohn, Liu, and Wardlaw (2022) was written, I instead referred authors to an ecology paper by O'Hara and Kotze (2010).

3.3 Discretization

Your main independent variable of interest is a dummy variable for whether a company is above median in X. This throws away a lot of information. Economically, it seems that the actual level of X will matter for Y, not just whether it is above or below median. Are the results robust to the full continuous level of X? Or, are the economic arguments for why the effect of X will be nonlinear, and the nonlinearities are such that the appropriate specification is to divide X into halves. I can see that a very high level of X might have an outsized effect on Y, but this would suggest using quartiles or quintiles.

There are a number of concerns with discretizing continuous variables – it throws away information, and it gives researchers the freedom to choose their quantiles – halves, terciles, quartiles, and so on. Discretization is justified if there are economic reasons for nonlinearities, but the onus is on the authors to provide the justification for the discretization, rather than just stating that they created a dummy variable. Similarly, some papers use a dummy variable for whether a company has at least one female director (rather than the number of female directors), whether a CEO has equity compensation (rather than the amount of equity compensation), or whether a disclosure in an annual report contains at least one word from a list (rather than the number or fraction of words from that list). Again, such a specification may be appropriate, but the authors need to motivate it. As with the prior section, no paper was rejected for this practice as it is fixable, but I sometimes made this comment as a suggestion.

4. Exposition

Papers were very rarely rejected solely due to exposition. However, when rejections were made on other grounds, I often provided expositional suggestions (as I sometimes did when issuing an R&R). They fall into three categories.

4.1 Clarity

I'm sorry to say that the paper is so poorly written that it would be difficult for a reviewer to provide an informative review and give constructive comments. I tripped up in virtually every sentence when trying to read your paper. For example, [examples of typos]. While none of these issues in isolation is major, added together they make the paper very difficult to read and any high-quality referee that I might send the paper to would either refuse to referee it or keep tripping up and thus fail to provide an informative review. In addition, they give the paper the overall impression of carelessness, and may lead to the referee wondering whether you have been equally careless in your coding / proofs.

It might seem obvious to proof-read a paper before submission, but many papers are submitted with elementary errors. It might seem that these issues should not affect the eventual publishability of the paper because they are fixable. However, as I later explain in Section 5.2, irremediability is not a necessary criterion for rejection – a paper may be rejected even if the issues are fixable, if it is too far from publication to converge speedily. Editors have limited time and there are trade-offs. If an editor has to use many rounds to get a paper to publishable standard, this will be at the expense of time spent on other papers to the detriment of the journal overall. Their objective function is the quality of the journal, not the quality of one specific paper.

Moreover, editors and referees are human, and what matters for their decisions and recommendations is not publishability but perceived publishability. A published finance paper should constitute (social) science, and if authors cannot even fix typos, then they may have been equally unscientific in other aspects of the research process. This applies not only to typos, but the appearance of tables, the consistency of bibliography formatting (e.g. having some journal names in italics and others in normal font), and other minutiae. When I was an analyst at Morgan Stanley, I was often asked to burn the midnight oil ensuring that line thicknesses were consistent between different graphs. As my boss explained, “if the client can't trust us to get the darn line thicknesses right, they won't trust us with a multi-billion-pound deal” (although he used a different adjective than “darn”). A similar principle applies to academic writing.

There is no economic significance in the abstract. One guideline (not a rule) is that the abstract of an empirical paper should contain one number of economic significance that the reader can take away, remember, and cite. For example, I can recite off the top of my head that Gompers, Ishii, and Metrick (2003) find 8.5% abnormal returns to governance-sorted portfolios, and that Holderness (2009) finds that 96% of firms have blockholders, as they both contain these numbers in the abstract.

This comment is self-explanatory. Similarly, there are many paper introductions that contain no, or very little, economic significance. While the goal of an abstract is to get the reader to remember the punchlines of your paper, the introduction should be a self-contained summary of the paper. Many referees have decided by the end of the introduction whether to reject or give an R&R to your paper, and then read the rest of the paper coming up with reasons to justify this initial hunch. Similarly, even if a paper is published, many readers will not read beyond the introduction. Just as a paper using instruments needs to explain them so the reader can assess their validity, any empirical paper needs to state the economic significance of key results so that the reader can evaluate their importance and plausibility. If the economic significance is very small, then the paper is uninteresting even if the results are statistically significant. If the economic significance is very large, its results may be implausible; for example, the authors may have failed to control for an omitted variable.

The sentence “A one standard deviation increase in X causes a change in performance by 2%” is vague. “Change” in performance is unclear, as it could refer to an increase or decrease. It is unclear how you are measuring performance – stock returns, operating profit over sales, ROA or something else. The timescale is also unclear – does performance change in the next year, after two years, or something else? How fast the performance change manifests is important for the plausibility of the results.”

Here, the introduction provided economic significance, but it is not clear what the economic significance referred to for three reasons. More broadly than statements about economic significance, “change” is often a vague word to use in a paper. Finance involves directional relationships, so it is clearer to refer to increases or decreases. The following are further examples where the writing could have been clearer, with the text in quotes and normal font, and my comments in italics.

“We show that board structure affects firm policies” (*what aspects of board structure? which policies? which direction were the effects?*)

“We use this conjecture to ascertain the consequence of addition or substitution of new equity components” (*consequence on which variables? which equity components?*)

“These firms consistently exhibit a different relationship between pay incentives and policy decisions.” (*which policy decisions? How is the relationship different – is it stronger or weaker?*)

“The association between incentives and firm financing decisions and increased focus turns significant” (*which financing decisions? what measure of focus? turns significant in what direction?*)

“When we split the sample, we find that ...” (*what variable are you splitting the sample by?*)

“The results are weaker in specification (5)” (*please describe specification in the text, e.g. “specification (5) when we add firm fixed effects”*)

The next extract is an example of claiming non-specific implications.

The sentence “Our result has several implications for executives, directors, investors, and regulators interested in sustainability” is vague. Any paper with “sustainability” in the title could include this sentence. It would be helpful to spell out what your specific implications are.

Many papers claim implications for practitioners and policymakers, to suggest that the findings are important. However, simply making the claim is unconvincing without giving specific examples of such takeaways. A sentence like the above often appears at the end of abstract. Rather than claiming several implications, and listing several constituencies to which it provides implications, it would be better to provide just one example.

The abstract of the paper should be labelled p1. This ensures that the page numbers on the PDF match the page numbers on the printed version, and makes it easier for editors and referees to provide comments. That way, if we refer to p9, it is unambiguous which page we are referring to.

This is a very minor comment that affects clarity, but is similarly very easy to fix. Some papers may have an unnumbered abstract page and start the introduction (which is the second page) with p1,

potentially to make the paper appear one page shorter, but this is an unhelpful practice. At the *RF* we introduced a box that authors have to tick when submitting to confirm that the abstract page is labelled p1, but even after doing so, papers are still mislabeled. This also applies outside the editorial process to papers that authors send to other researchers asking for comments; consistent labelling will help them do so.

4.2 Length

Your introduction is 11 pages, which is far in the right tail of introductions that I receive, most of which max out at 6 pages. This makes it hard to for the reader to see the paper's punchline, or s/he will run out of steam before getting to the actual analysis. For example, you spend 3 pages (p2-4) setting the scene and describing the related literature before you actually get to your paper. Then you finally explain your research question at the top of p5 but flip back to the related literature at the bottom of p5. Throughout the introduction, it goes back-and-forth between your contribution and related literature like a Hamlet soliloquy. A better approach would be to start with a single paragraph describing the context, and then to go immediately to your analysis – your hypotheses, identification strategy, data sources, and results. Then, and only then, you can discuss the related literature and how you differ. It is essential to acknowledge related papers, but it is difficult for the reader to evaluate your contribution relative to them without first knowing what your contribution is.

As stated in the prior section, the introduction is extremely important as many referees will have made their decision, or at least formed a strong opinion, by the end of it. Thus, it needs to be as crisp and punchy as possible, describing all the essential details of the paper. For an empirical paper, it should contain the hypotheses, identification strategy, how the main variables are measured, and the economic significance. For a theoretical paper, it should sketch the model – who the key players are, what actions they take, and what information they have – and describe the primary intuitions behind the results.

Equally importantly, the introduction should not contain more than the essential details of the paper, just as a witness in a trial needs to tell both the whole truth and nothing but the truth. One of my most common phrases in decision letters is how expensive “real estate” is in the introduction. It

is even pricier in the first page and dearest of all in the abstract. Thus, that a sentence is relevant and adds value is not sufficient to include it; it must add enough value to justify the cost.

The paper spends far too long explaining the importance of innovation. It contains two pages explaining that innovation is important, plus two figures which – in the published version – would appear in the introduction where “real estate” is expensive. But any reader of an academic journal already knows this.

This extract illustrates a common reason for excessive length – a superfluous motivation of the paper. Introductions highlighting the importance of climate change or how devastating COVID-19 was often fall into this category. Moreover, these introductions focused on the importance of the context, rather than the specific question the paper is asking within this context.

The use of footnotes could be more measured. Having lots of footnotes gives the impression that the authors are not really sure what is central to the paper and are hedging. They also impose lots of burden on the reader as a good referee needs to read all of them to ensure that he/she does not end up making a comment that is addressed in a footnote. It also causes the paper to have a staccato feel as the reader needs to flip back-and-forth constantly between the main text and the footnotes. A good guide (although not a strict rule) is no more than one footnote per page on average. I would encourage you to think very carefully about what is truly peripheral (and thus can be deleted entirely) and what is central and should be promoted to the main text. Examples follow below, but I encourage you to go through the whole paper and verify whether each footnote is necessary:

- a. You have six footnotes on p2 to provide evidence that [topic of paper] is important. None of these seem necessary – any RF reader will know that [topic] has generated interest.*
- b. Footnote 7 is unnecessary. If the reader wants more detail on [a minor issue], she can Google. It’s not clear why the link you include is more informative than other links available on the web.*

Footnotes and internet appendices should not be seen as freely disposable. Footnotes take up space in the main paper, and internet appendices need to be referred to in the main paper. Plus, they require a careful referee to read them.

Another reason for excessive length is “nice-to-have” extensions that neither test an important hypothesis nor address an important methodological concern. Even if authors think that some readers

might find an extension interesting, so that it has strictly positive value, the bar for inclusion should not be zero. Applying the social welfare criterion, the value of an extension should exceed not zero, but the cost of reading it. Moreover, referees may form an assessment based on the average quality of a paper – even though three major results and one minor one should dominate three major results alone, the latter may lead to a more positive overall impression. Thus, a potential alternative bar for including an analysis is whether it “raises the average” of a paper.

4.3 Citations of Irrelevant Papers

Your paper has a seven-page reference list, which is far in the right-tail of bibliographies that I see. At times, it reads like a Masters thesis, where you try to cite as many papers as possible to show off your knowledge of the literature, when those papers are only tangentially relevant. This is in particularly the case in the introduction, where “real estate” is particularly valuable and you should be focusing on your unique contribution. The constant citations of unrelated papers interrupts the flow and distracts the reader from your paper.

After your 7-page introduction, you then have another 5 pages of literature review (which is typically included at the end of the introduction), which leads to an effective 12-page introduction.

Both of the above extracts describe a fundamental trend in recent years: the growing length of bibliographies. This is partially justified as the volume of research has risen in recent years, and thus the number of related papers (which indeed must be cited) has increased. However, many of these citations are to irrelevant papers. This is a problem for two reasons. First, it distorts the paper and distracts the reader from the paper’s own contribution. Second, it distorts the profession. Rightly or wrongly, citations are an increasingly important measure of a researcher’s quality and impact. Under current citation practices, a researcher can increase her citation count by working in a popular field, even if she makes only a minor contribution. When I was an assistant professor, one of my senior coauthors added a couple of citations to our paper, saying that it’s nice to cite friends and colleagues. It may be nice, but it’s not scientific; just like it’s nice but unscientific to accept a paper by someone you like.

Citations also affect a journal's impact factor and thus reputation. Journals can similarly increase their impact factors by publishing papers that make marginal contributions on a popular topic, or show a previously-documented result in a new and topical setting. For example, if X has been shown to affect Y in prior economic downturns, a paper demonstrating this result during the COVID-19 pandemic would likely be highly cited because COVID is a topical issue, even if it doesn't cause the typical reader to Bayesian update. While attention has rightly been paid to rule out errors of omission (failure to cite relevant papers), there is little attention to reduce errors of commission (citing irrelevant papers), even though the latter skews authors' and journals' reputations and incentives.

While the above extracts illustrate the general practice of excessive citations, the following extracts will give specific examples. I will predominantly use examples where I asked authors not to cite my own papers, so that I am not misconstrued as suggesting that other papers should be cited less. However, I will use some other examples where I have no paper that makes that specific point, with apologies to the authors. I will give a sample citation in quotes and normal font, and then the extract from my decision letter in italics. The sample citation is paraphrased to avoid identifying the manuscript; P refers to a paper being cited.

"The importance of ESG has risen over time (Edmans, 2023)." *That ESG has become more important over time is institutional fact, and not a contribution of Edmans (2023). That paper uses the importance of ESG over time as motivation, but it did not discover the increasing importance of ESG.*

"The board's objective function is shareholder value, since directors' fiduciary duties are to shareholders (P)."
That directors have fiduciary duties to shareholders is institutional fact, and not a contribution of P . P merely used this fact to motivate an assumption in their model.

"In equation (5), K refers to intangible assets such as employee satisfaction (Edmans, 2011)."
That employee satisfaction is an intangible asset is institutional fact, and not a contribution of Edmans (2011). It used the fact that employee satisfaction is an intangible asset to motivate the analysis of whether it is fully valued by the market. However, in your model, K is fully priced in, so you do not use this finding and no citation is necessary.

"Total equity is made up of prior equity holdings and newly-vested equity (Edmans, Fang, and Lewellen, 2017)."
That equity can be either previously-vested or newly-vested is institutional fact,

and not a contribution of Edmans, Fang, and Lewellen (2017). They used vesting equity as an instrument for equity sales, but did not discover the idea that equity vests.

“ $e = e_L$ corresponds to shirking or enjoying the quiet life (Bertrand and Mullainathan, 2003).” That managers may enjoy the quiet life was already known long before Bertrand and Mullainathan (2003). That is a seminal paper, but their contribution was to show that weak corporate governance causes managers to shirk, not to point out that managers may shirk.

The first three extracts are examples of a similar point – papers are often cited for institutional facts that they contain or use, often to motivate the context (extract #1), a modelling assumption (#2), an empirical hypothesis (#3), or an identification strategy (#4), even though these institutional facts were known long before those papers were written. The paper did not introduce these facts; even if those papers had not been written, the manuscript could have still stated the fact. Thus, these papers receive citations for working on a popular topic, not for contributing a result. The fifth cites a paper for a term that it uses, rather than a result that it finds. (In contrast, if a paper coins a technical term, then a citation may be warranted.)

The following extracts from decision letters should be sufficiently clear that no quote from the paper is necessary.

I'm not sure you should attribute Tobin's Q to [list of papers]. Tobin's Q should be attributed to Tobin. Literally thousands of people calculate Tobin's Q in this way; it is not clear why you are singling out those papers.

I'm not sure you should credit [list of papers] with using control variables. Any empirical corporate finance paper controls for control variables. You can simply say that you control for other variables that may be correlated with Y . If prior papers have shown that the specific control variables that you use affect Y , you should cite them, but you should not cite papers for the general idea of using control variables to address omitted variables bias.

I'm not sure you should credit [list of papers] for fixed effects. Any empirical corporate finance paper controls for fixed effects to address time-invariant unobservables.

I'm not sure you should credit [list of papers] with highlighting that there may be reverse causality. Any researcher in corporate finance knows about reverse causality and tries to address it; even had those papers not been written, it should be obvious to any reader that there are reverse causality issues in your setting. Of course, if other papers have explicitly documented the reverse

causality, or introduced instruments that address reverse causality, you should cite them, but you should not cite papers for the general idea that reverse causality is a concern.

Lots of papers used the MSCI KLD database before the paper you cite on p25, so it's not clear why you single out this one.

In the above examples, papers are given credit for methodologies, variables, or datasets that they did not discover, because they were well-known before those papers were written. Often, those papers are cited because they use the methodology (e.g. fixed effects) in a similar context to the manuscript, which leads to papers being cited for being in popular areas rather than making significant contributions. Even if no prior paper used fixed effects in that precise context, the authors of the manuscript would have known to control for fixed effects. Sometimes, a paper will say “We use the MSCI KLD database as it is the most commonly-used measure of ESG”, and cite five papers that use this database. However, citing five papers is not evidence that MSCI KLD is the most common measure; a researcher could equally cite five papers that use Thomson Reuters (an alternative dataset). The paper can simply make the statement with no citations, or citing an authoritative survey paper which makes this point.

Other times, papers are mis-cited – they are credited for topics that they do not actually study.

The following are examples:

Edmans (2009) is about blockholders, not family ownership, and is not about risk exposure or risk-taking.

Porter (1992) is not a formal model of short-termism. In [myopia model], there are no outside investors; there is just a market (there is no measure of outside investor pressure). That disclosure alleviates information asymmetries is automatic, and not a contribution of [paper].

The first is an instance where the authors cite a paper that covers a general topic, when their own manuscript it is on something much more specific. The second concerns three issues: a reference to a conjecture rather than a research finding, a model that does not contain the element that is cited for,

and another theory that did not contribute a particular result – instead, it used it as an input into the model.

Some references to prior research are rather loose. For example, on p8 you refer to very many papers in the first full paragraph, and then makes a vague claim that “we introduce a novel channel connecting both fields of research”. However, it’s not clear how you “connect” these fields of research, or that they even need to (or should) be connected. For example, even if X affected firm value, it might not affect future stock performance if it were fully priced in.

You write “our paper is related to the broader literature on corporate governance, such as boards of directors”. Your paper has a clear and focused research question on CEO pay. It is unrelated to the literature on boards of directors. They have no implications for your research, nor does yours for theirs. Even had those papers not been written, your research design and implications would be unchanged.

While the prior examples were of superfluous citations of specific papers, these examples concern superfluous citations of entire literatures. Sometimes papers will claim connections to several different literatures to give the impression they have broad impact, and thus merit publication in a general interest journal. However, the connections to the literatures are vague and non-specific. Sometimes, a manuscript will say “we have implications for the literature on dividend policy”, and cite many papers on dividend policy, without actually saying what those implications are. Even if the manuscript had implications for dividend policy, it could cite a survey paper on the topic, rather than lots of individual papers. Which papers authors choose to cite and ignore is often arbitrary, or sometimes strategic with papers by more famous authors being more likely to be cited.

The second example concerns a paper on one area of corporate governance, and cited an entirely separate area of corporate governance that it bore no relation to. A branch of a tree has no direct connection to another branch of the same tree.

I was unsure why you try to highlight that ESG matters more during [type of corporate event] and [specific institutional context], except for trying to strategically cite these papers. You do not focus on these contexts. If anything, these citations reduce the relevance of your paper, as they suggest that ESG does not matter much outside these contexts.

The excerpt is self-explanatory.

5. General

5.1 Advice to Authors

Despite the negative recommendation, the referee endeavors to be constructive and makes several suggestions for additional analyses and robustness checks that you could conduct. However, my advice is not to follow any of these suggestions. The only reason for rejection is that the contribution is unfortunately not strong enough for a top general-interest journal. The additional analyses are unlikely to change the magnitude of the contribution, and you have already shown that your results are robust. My advice is not to revise the paper substantially, but simply to send the paper with minor changes to a journal with a slightly lower bar.

I sympathize with the fact that, when submitting to a journal, authors hope that, even if they get rejected, they get lots of suggestions for things to do so they feel the rejection has been “productive” and they can do something in response. However, I think the referee has done you a great service here by providing a short report focused on the contribution. A long list of comments is unlikely to help you, as you may end up spending many months incorporating them when they are unlikely to improve the paper – the paper is already well-executed, and the main reason for rejection is contribution. I believe your optimal strategy is not to revise the paper substantially, but simply to send the paper with minor changes to a journal with a slightly lower bar. I rarely write this comment for papers that I reject because they are sometimes poorly written and executed, but this is not the case here.

Sometimes, authors are upset with short reports, but a short report can often be the most useful. In contrast, referees might be afraid to send a short report, and thus pad it with suggestions. The “social welfare” criterion mentioned earlier is relevant here. It might take the authors many months to implement the suggestions, and if doing so has a negligible effect on the paper’s publishability, the referee’s comments destroy social welfare.

The goal of a referee report is to provide a recommendation to the editor whether to reject the paper or invite a revision. It is not to improve the paper for the next journal – authors should instead obtain such comments by presenting at seminars and conferences and sending their work to colleagues and researchers in the same field. Of course, when reading the paper to assess its publishability, the referee may find comments to improve the paper, and it is good citizenship to provide those comments

even if they did not ultimately affect the decision. However, if the paper is a clear reject based on the contribution, there is no need to force yourself to provide suggestions on the execution.

Similarly, the role of an editor is to provide a fair and objective decision to the authors, and to faithfully communicate the reasons behind it. It is not to improve the paper for the next journal, or to convince the authors that she spent several hours reading the paper. If the reason for rejection is that the editor has received an excellent report from a trusted referee, the editor can communicate this to the authors and move on. She should not feel obliged to pad out the letter by repeating the referee's comments if they are already clear, nor to try to come up with suggestions of her own, which may destroy social welfare if the paper is outside her field.

Please don't hesitate to ask if some of my comments are unclear, or to say if you disagree with them. This is not a perfunctory offer, as sometimes on my own R&Rs, my authors and I have spent countless days debating on what the editor actually meant by a comment. Or, we think that an editor's comment is wrong but we don't think we can tell simply him that he's wrong and ignore the comment; instead, we spend ages thinking of a new unnecessary analysis to avoid telling him he's wrong but to pretend we are "addressing" his concern. If you think that some of my comments are mistaken, feel free to explain to me why and ask whether you should still address the comment in the light of your explanation. If am convinced, I will withdraw that comment so you don't need to do anything. Also, if I can do anything else to help speed up the revision (e.g. if a Word version of this letter would help you copy-paste my comments into the response document), please say.

I often included the above paragraph in R&R letters. Social welfare is destroyed if the editor or referees have made a nonsensical comment and the authors feel pressured into conducting an analysis to address it, or if a comment is unclear and authors waste time trying to understand it. While different editors may have different views on this, authors should probably feel more comfortable reaching out to editors between rounds than is currently the norm, even if the editor has not given a specific invitation to do so.

My official decision is "reject-and-resubmit". In the rest of this letter, I explain what this decision means, and what the conditions under which you should choose to resubmit.

What do I mean by a "Reject-and-Resubmit"?

I have sometimes received this decision myself, and am confused as to what the difference is from a weak revise-and-resubmit. I mean the following by this decision. A “weak revise-and-resubmit” indicates that the referees’ comments are difficult to address, but if the authors are somehow able to address, likely the paper proceeds to the next round. There is an implicit understanding that, if all the concerns are addressed, the paper will move forward. Unfortunately, this is not the case of for a “reject-and-resubmit” – even if you do address the main concerns of the referees, the paper may still end up being rejected. Roughly speaking, a “revise-and-resubmit” is evaluated more on its incremental changes versus the original submission, and whether they have satisfactorily addressed the reviewers’ concerns – the research question has already been deemed to have sufficient potential, which is why an R&R was given. A “reject-and-resubmit” is evaluated more in its entirety, as a brand new submission, including whether the research question is sufficiently interesting.

Think of it this way. Assume that the paper was given a straight rejection, but you believe that the referee’s comments are good and you revise the paper to address them. You would normally then think about what journal to send the revised paper to next. With a “reject-and-resubmit”, you may include the Review of Finance within that choice set. If you sent the paper to a new journal, there would be no guarantee that the paper will be accepted even if you have addressed all of the RF referees’ comments. Similarly, if you send the paper back to the RF, there is no guarantee that it will be accepted, even if you have addressed all comments. But, I am at least willing to give you the option to try.

While it might seem that a “weak R&R” would be preferable to a “reject-and-resubmit”, I do not want to lead you on. I recognize that revising a paper is a monumental task. It might take you several months to address all the referee’s empirical and expository comments, and the paper might still not be publishable if the research question that your new paper is centered around is not sufficiently strong. Thus, this is not simply a question of “responding to the referee’s comments” as with an R&R, but finding a topic for a new paper. On the other hand, it also means that you don’t need to respond to every single referee comment. He suggests a potential topic for a new paper, but you might think of an alternative question.

What are the conditions for likely success?

[Paper-specific comments]

What should you do?

Of course, it is your decision as to what to do with the paper going forwards. Here is my personal advice, which you can obviously take or leave.

First, decide whether you think the referee’s comments will indeed improve the paper – i.e., even if you had no option to resubmit to the RF, would you incorporate them before submitting to the next journal? If not, then it is probably best to move on. Since the chance at the RF is low, you should not spend several months revising the paper for the RF reviewers if you don’t think that doing so would help you at other journals.

If so, then please revise the paper to address the comments, and prepare a point-by-point response document to both referees (as if you had an R&R). If you do choose to resubmit the paper, please

note that – in the interest of streamlining the process – I will desk-reject the paper without sending back to the referee if I feel that you have not gone sufficiently far in addressing his comments for the odds of success to be realistic.

While some journals no longer allow for reject-and-resubmit decisions, the *RF* and other journals still do. Authors are often confused about what this decision means, and whether they should bother trying to revise the paper or instead bite the bullet and send to another journal. The above is an extract from a reject-and-resubmit letter giving the authors advice on this decision, with the acknowledgement that different journals may mean different things by a reject-and-resubmit.

5.2 Reaching an Editorial Decision

These are comments where I explain how I weighed up the referee reports and my own reading to reach a decision. These comments are shared in case they help authors understand how editors reach decisions (with the acknowledgement that different editors may use different criteria).

I recognize that endogeneity is sometimes a “cheap shot” that referees use to reject the paper (sometimes claiming that a variable is endogenous without explaining why it might be endogenous, or why endogeneity is a fatal concern for the paper), but this is not one of these cases. Starting with the former, the referee explains multiple potential sources of endogeneity in his report. Turning to the latter, I am not a member of the identification police and believe that many papers can make first-order contributions by simply documenting interesting correlations or new facts, even if they cannot make causal statements. Unfortunately, the specific research question of your particular paper does require you to be able to make causal statements.

Our profession’s increasing focus on identification is a strong positive step. Papers documenting only correlations should not claim causation, and indeed this focus is an important differentiator of academia compared to consultancies that are increasingly pumping out research with causal claims. However, almost all trends have both costs and benefits. This focus on identification sometimes allows referees to recommend rejection by crying endogeneity. However, a mere allegation of endogeneity should not lead to rejection; the referee needs to point out specific omitted variables, and

explain why they would bias the result in the authors' favor rather than against them. Similarly, it may be that the goal of the paper is not to show causality but only to document correlations. Thus, a mere mention of endogeneity does not automatically lead to rejection; it needs to be that the concerns are valid and also significantly weaken the paper's contribution.

More broadly, I understand that some of the most important questions cannot be identified with absolute precision. If a paper's contribution is first-order, I might be willing to cut the authors a little slack on the identification (as long as they are up-front that they do not show causality). However, here the referee has significant concerns on the contribution as well as the identification, which makes it difficult to move forward.

Almost no paper is perfect. The bar for publication is that a paper should make a substantial contribution to knowledge, not be perfect. Thus, a paper may be publishable if it is imperfectly identified but on an important topic, or in a crowded literature but precisely identified. Often papers may be rejected despite neither the identification being non-existent nor the contribution being zero, but because neither is sufficiently strong to give grounds to publish the paper.

More generally, it is not that the default decision is to accept the paper and that, only if you find something wrong with it, should the paper be rejected. Instead, the default decision is to reject the paper; only if you find something new, interesting, and important in it, should it be published. Sometimes, authors are disappointed to be rejected because they think that most of the referee's concerns are fixable, so there is little wrong with the paper – but this is insufficient for publication. A paper is not accepted if it has “little wrong”; it has to have “a lot right”. However, referees and editors should also take responsibility. Some referee reports are what a former editor described to me as “sniping” – the referee doesn't find the paper interesting, but doesn't have the confidence to say so and instead gives a long list of small to medium-sized concerns about the paper, even if none is pivotal or irremediable. Instead, a rejection report should highlight either why the problems are unfixable, or whether the paper remains uninteresting even if the problems could be remedied.

Similarly, editors should not simply point out flaws (or hide behind a referee report pointing out flaws) but explain why the contribution is insufficient even if the flaws were fixed.

At the RF, we have a second round up-or-out policy to avoid papers being dragged through many rounds. Unfortunately, here the referee has many serious concerns which mean that the paper is a long way from publication. It is very unlikely that it will converge within one round and thus, regrettably, the prudent decision is to reject the paper. I recognize that an R&R might give authors the best short-term news, but it may well be damaging in the long-run since the authors spend several months revising the paper and get rejected at the second round because, despite significant improvement, the paper is still unpublishable.

A paper may be rejected even if the concerns are eventually fixable, if fixing them will consume substantial journal resources. A rejection need not mean that the paper can never be published, but that it is too far from publication to converge rapidly. (Otherwise, taken to the extreme, authors could provide a promising research proposal and expect the editor and referees to suggest the identification strategy). Obviously, this means that authors should make the paper as strong as possible before submitting it to a journal, rather than relying on editors and referees to improve it.

Now I fully understand that theory papers should not always be motivated by trying to explain empirical facts. I have long believed that the profession judges theory papers too much on their outputs (can they fit the data) rather than their inputs (are they modelling interesting economic phenomena). The former is problematic because (a) you can always reverse-engineer a model to fit data; (b) it may be that the real world is not optimizing – e.g. Holmstrom's model of relative performance evaluation might be rejected if it were written today since RPE is rare in reality; (c) it may be that you're introducing a new theory and the empirical support will come later. However, to justify publication of a model based on its inputs, it's necessary to be convinced that these inputs are capturing first-order empirical phenomena. For Holmstrom, it is reasonable to think that output is affected by noise and that agents are risk-averse so would like this noise to be filtered out, but it is unclear that your new input is a first-order concern in the real world.

The above extract is from a decision letter on a theory paper. Even if a theory paper's predictions are not supported by the empirics, this should not be grounds for rejection. It may be that the real world is not optimizing (in which case the authors should be up-front that their model is normative, not positive), or that there is not yet a sufficiently accurate dataset or clean identification strategy to

test the predictions. If this is the case, the model should, in part, be judged on the inputs – whether the new forces that the model is capturing are first-order. While importance is partially subjective, empirical results (using either archival data or survey responses) that document the importance of these forces is often helpful.

5.3 Emails Between Authors and Editors

I'm afraid that, as stated at <http://revfin.org/aims-and-scope/>, we are unable to provide an indication of whether a paper is suitable for the RF before actual submission. This is because providing any such guidance would be irresponsible without reading it carefully (and, in some cases, getting external input) - else we might end up saying that a paper is unsuitable when, under a close reading, its errors might be fixable. Alternatively, we end up saying a paper is suitable and, no matter how carefully we caveat this and stress there are no guarantees, authors get upset when we reject them because they feel we have led them on. In terms of the topic, we are unable to provide further guidance than it has to make a major academic contribution to a broad finance audience and we have the same standard as the top-three journals – it wouldn't be responsible to give a decision based on the broad topic category without understanding the paper's specific contribution to the topic, what other contributions have been made etc.

Sometimes authors emailed me (or even sent me through LinkedIn Messenger, which I do not recommend) a manuscript and asked whether it would be of interest to the RF. Occasionally the authors would couch it as asking whether the paper met the “scope” of the journal rather than enquiring about the quality. I was unable to make either type of assessment, and the above extract explains why.

I'm afraid that our policy (as with any journal) is that a rejected paper cannot be resubmitted, even if the authors believe that they can address the referee's concerns. Virtually all authors think that they can address the referee's concerns, and if we opened the door to that, then journals would never be able to go about their work, as many rejected authors would want to resubmit their papers. In addition, while a referee report recommending R&R is supposed to be an inclusive list of concerns (so that if they are addressed, the referee should accept on the second round), a rejection letter does not serve the same purpose – the referee may recommend rejection without necessarily listing all the issues he may have with the paper. The list could likely be incomplete but sufficient to make the recommendation; it is not the quasi-binding contract that an R&R report should be.

Rejected authors sometimes emailed me saying that they believe (often saying “strongly believe”) that they can address the referee’s concerns and asking if they could resubmit the paper. The above extract explains why not. In addition, as explained in Section 5.2, a paper may still be rejected even if the concerns are fixable, if the contribution is not sufficiently strong or if the paper is so far away from publication that it is unlikely to converge rapidly. Of course, if rejected authors believe the referee’s concerns were incorrect (rather than simply being addressable), they could pursue the appeals process.

Please note that the short time between getting the confirmation of editorial assignment and this decision email doesn’t mean that I only spent a few minutes on the paper; I read all papers before deciding which editor to assign.

I desk-rejected a paper if I was able to assess the magnitude of its contribution or quality of its execution myself without the help of referees. As Managing Editor, I saw all papers before they were assigned to editors. I would read them to decide whether to desk-reject them, and if not, which editor to assign them to.⁴ If I decided to desk-reject the paper, I would assign the paper to myself and then write the desk-rejection letter. The way that Editorial Express (the *RF*’s online system) works is that, once the editor is assigned, the authors receive an email notifying them that the paper is under review and who the editor is. As a result, the desk-rejection letter might follow as soon as 15 minutes after the assignment email. This sometimes led to unhappy emails from authors who thought that I had spent 15 minutes reading the paper and writing the rejection letter. (I would have probably thought this myself had I not known how the assignment process works.) Instead, I had already read the paper before Editorial Express sent the email notifying the authors that I was the assigned editor. I often

⁴ The editor could subsequently decide to desk-reject the paper herself; this might be particularly the case if the paper was outside my field and so I could not assess its contribution.

included the above excerpt in reply to author concerns, or in decision letters to pre-empt such concerns.

5.4 Emails Between Editors and Referees

Shortly after I took over as Managing Editor, one senior person in the profession advised me to provide feedback to referees – of all seniority – on how they can make their reports even better. Might I take the liberty of providing some tips here? While I think you pointed out some important issues, I think your report could have been a bit more informative in tackling the central issues of the paper – many of your comments, while insightful, were on more ancillary parts of the paper. It would also be useful to have a more detailed cover letter, on the pros and cons of the paper so that I will know what to weigh when making my decision. Some parts of the report were difficult for me to understand without reading the paper (e.g. what the A and B hypotheses were). Some feedback I provided to the authors is below, in case it's helpful. I am not saying that all my feedback is right; more that these are the more central issues for the authors to address.

Referees are extremely important to the editorial process, and thus the profession more generally, but receive almost no constructive feedback on the quality of their reports. PhD programs teach students how to write papers, but not how to write referee reports (beyond perhaps an assignment in a PhD class), and there is no training after the PhD. Sometimes editors will provide positive feedback when thanking referees, but negative feedback is very rare. I sometimes tried to highlight issues that referees missed, or advise them that they focused too much on the minutiae. However, I provided far fewer such emails than I might have done if this were the norm, and was likely too timid in doing so. It is the norm for discussants to give constructive feedback on papers at conferences and so there is no embarrassment for authors to receive them; in turn, this makes discussants more comfortable giving such feedback. A similar norm would be useful for referee reports. Of course, providing such feedback takes time, but should help editors in the long-run through improving refereeing quality.

6. Conclusion

This article has provided extracts from decision letters, and emails to authors and referees, from my tenure at the *Review of Finance*. The intent is to help authors improve their papers and to increase transparency on how editorial decisions are made and how the editorial process works. This is only a first step, and only contains extracts from a single editor from a single journal over a single time period. It may be that other editors have different views that they are willing to share.

Similarly, there might be other resources created through the editorial process that may be worth sharing. Just as we receive very little training on how to write referee reports, we receive very little training on how to respond to referees. To throw out a blue sky idea, just as journals have websites that contain online appendices, they might also contain the referee reports and response documents associated with published papers, if the referees and authors agree to share them. In addition to providing insights on how to write and respond to referee reports, many analyses are conducted during the revision process that never see the light of day, even in an online appendix – such as additional cross-sectional tests that were insignificant. Since the referee thought that such additional analyses were useful, a separate set of authors might think that they are sufficiently interesting to center a new paper around. Knowing that the results are insignificant will save them the effort.

Negative referee reports are also informative, particularly because referees often spend more time on papers than editors. This article has focused on editorial letters because I am not the owner of referee reports and thus not at liberty to share them, but an even more blue sky idea could be a repository for referee reports – if an editor believes that a report has been particularly useful, she can ask the referee if he is willing to contribute it anonymously to the repository. In addition to helping others learn how to referee, it also helps the profession understand the contribution of a particular paper. At present, authors can engage in referee shopping, sending a manuscript from journal to journal hoping for an inattentive referee. In some cases, a referee I assigned informed me he had

already rejected the manuscript for another journal based on insufficient contribution relative to an uncited paper, yet the authors still did not cite the paper when submitting to the *RF* – hoping to find a referee unaware of this paper. It would also deter authors from submitting manuscripts to journals prematurely, hoping that the referees will help improve them. It would reduce the loss to social welfare arising from an expert referee spending a day providing insightful comments, for the authors to shop the paper to another journal and take up a day of another referee’s time to provide very similar feedback.

However, the negative consequences of such a practice would need to be considered very seriously. An uninformative and unfair referee report may prejudice referees at other journals, but this might be mitigated by editors only contributing top-quality reports to a repository. Authors may be unwilling to submit to a journal if they know that negative referee reports will become public. However, authors are willing to submit papers to conferences even though the discussion is public, and authors are particularly pleased when accepted to prestigious conferences even though this means that the discussion is more likely to be seen by future referees. Some discussants post discussions on their website, without seeking the authors’ permission, which has a similar effect to a repository for referee reports. This helps the profession understand the paper, but may prejudice future referees, particularly since conferences (unlike journals) encourage submission of early work.

The editorial process consumes a substantial amount of inputs, but in doing so it produces a substantial amount of outputs, above and beyond the published papers. This article seeks to take a first step in sharing some of these outputs more broadly, focusing on editorial decision letters, and to initiate a discussion on whether other outputs might also be shared.

References

- Allen, Franklin and Alex Edmans (2017): “Editorial.” *Review of Finance* 21, 1–6.
- Berk, Jonathan B., Campbell R. Harvey, and David Hirshleifer (2017): “How to Write an Effective Referee Report and Improve the Scientific Review Process.” *Journal of Economic Perspectives* 31, 231–244.
- Bertrand, Marianne and Sendhil Mullainathan (2003): “Enjoying the Quiet Life? Corporate Governance and Managerial Preferences.” *Journal of Political Economy* 111, 1043–1075.
- Butler, Alexander W. and Timothy Falcon Crack (2022): “A rookie’s guide to the academic job market in finance: The labor market for lemons.” *The Financial Review* 57, 775–591.
- Cochrane, John H. (2005): “Writing Tips for Ph.D Students.” Working paper, Stanford University.
- Cohn, Jonathan, Zack Liu, and Malcolm Wardlaw (2022): “Count (and Count-Like) Data in Finance.” *Journal of Financial Economics* 146, 529–551.
- Edmans, Alex (2009): “Blockholder Trading, Market Efficiency, and Managerial Myopia.” *Journal of Finance* 64, 2481–2513.
- Edmans, Alex (2011): “Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices.” *Journal of Financial Economics* 101, 621–640.
- Edmans, Alex (2022): “The Purpose of a Finance Professor.” *Financial Management* 51, 3–26.
- Edmans, Alex (2023): “The End of ESG.” *Financial Management* 52, 3–17.
- Edmans, Alex, Vivian W. Fang, and Katharina A. Lewellen (2017): “Equity Vesting and Investment.” *Review of Financial Studies* 30, 2229–2271.
- Gompers, Paul, Joy Ishii, and Andrew Metrick (2003): “Corporate Governance and Equity Prices.” *Quarterly Journal of Economics* 118, 107–155.
- Gormley, Todd A., and David A. Matsa (2014): “Common Errors: How to (and Not to) Control for Unobserved Heterogeneity.” *Review of Financial Studies* 27, 617–661.
- Harvey, Campbell R. (2017): “The Scientific Outlook in Financial Economics.” *Journal of Finance* 72, 1399–1440.
- Hirshleifer, David (2015): “Cosmetic Surgery in the Academic Review Process.” *Review of Financial Studies* 28, 637–649.
- Holderness, Clifford G. (2009): “The Myth of Diffuse Ownership in the United States.” *Review of Financial Studies* 22, 1377–1408.

- O'Hara, Robert B. and D. Johan Kotze (2010): "Do Not Log-Transform Count Data." *Methods in Ecology and Evolution* 1, 118–122.
- Pedersen, Lasse Heje (2021): "How to Succeed in Academia or Have Fun Trying." Working Paper, Copenhagen Business School.
- Porter, Michael J. (1992): "Capital Disadvantage: America's Failing Capital Investment System." *Harvard Business Review* 70, 65–82.
- Roberts, Michael R. and Toni M. Whited (2013): "Endogeneity in Empirical Corporate Finance." *Handbook of the Economics of Finance*, volume 2, ed. George Constantinides, Rene Stulz, and Milton Harris. Elsevier, Amsterdam.
- Spiegel, Matthew (2012): "Reviewing Less – Progressing More." *Review of Financial Studies* 25, 1331–1338.
- Weisbach, Michael S. (2021): *The Economist's Craft: An Introduction to Research, Publishing, and Professional Development*. Princeton University Press, Princeton.
- Yermack, David (2017): "Corporate Governance and Blockchains." *Review of Finance* 21, 7–31.