Some Research and Writing Tips

Part 1: Research

Ben Jacobsen
University of Edinburgh
&
New Zealand Institute of Advanced Study

ben.jacobsen@ed.ac.uk

This Draft: 6 March 2014

ALL COMMENTS ARE WELCOME

This guide gives tips on how to do (quantitative) research and on how to write it up. These insights seem to hold throughout all (social) sciences. I illustrate them using examples from finance and economics. My main goal is to save budding researchers time by preventing them from doing bad research and helping them to attract attention to their research by writing it better. My insights are based on learning from the best and from many of my own time-wasting mistakes over the last twenty-five years.

JEL-codes: A1, B1, C1, D1, E1, F1, G1, H1, I1, J1, L1, M1, N1, O1, P1, Q1, R1, Y1, Z1

Acknowledgements: I thank Utpal Bhattacharya, Jo Danbolt, Paul Geertsema, Jens Hagendorff, Susan Hancock, Ufuk Gucbilmez, Ben Marshall, Bill Rees and Paolo Quattrone for helpful comments and suggestions.

Disclaimer: These are my views. Feel free to prove me wrong and get top publications your own way.

In my experience, (quantitative) research[1] in all fields is essentially based on the same underlying principles. Unfortunately, guidelines regarding these principles are rare. Most students have to learn it the hard way by picking up bits and pieces of advice along their academic way.[2] Most researchers not only lack advice on doing research, but they also get precious little advice on how to write up their findings properly.

In this guide[3], I give my personal views on what separates good research from bad research (Part 1) and what separates good and bad writing (Part 2). My hope is that it prevents budding researchers from wasting time on research that was doomed to fail from the start - like some of my own - and helps them to achieve a higher level of research quality (and publications). I will rely on examples in finance and economics, often related to my own work simply because I know the details and the reasoning behind the choices I made and I prefer to criticize my own work rather than the work of others. Last but not least, these examples are also easy to explain to a non finance or a non economics audience.

This guide is intended for researchers like myself who can use a little help every now and then. Of course, what is true in the Arts also holds in the Sciences: the truly gifted can flaunt any guideline and still research and write brilliantly.

A good research question 1) is simple and 2) makes a large contribution to the literature. Let's start with 'simple'.

1.1.1 Use unambiguous, well-defined concepts

Consider this research question: "Are stock market returns higher in January?"

[1] While many principle I discuss here also apply to qualitative research my foces is on quantative research.

[2] Surprisingly many of the 'how to write and do research' guides also beat around the bush and fail to address what I would consider the 'main issues' and do not go much further than suggestions that an abstract should summarize your research.

[3] This guide is based on: my own research and writing experience, my teaching of 'How to do research' classes for Master and PhD students, workshops taught for staff and PhD students in universities around the world, my experience from research projects with staff and PhD students, and my experiences as an editor, associate editor and a referee. Most of all, I have been fortunate enough to learn from some of the best, and less fortunate in that I have made many of the mistakes I discuss here myself.

The first part is easy: ‘stock market returns’ is a relatively well-defined concept. If you use major stock market indices from a number of different countries, like the S&P 500 or the FTSE 100, people will accept your ‘concept’. ‘January’ is also a well-defined concept and so is ‘higher’. ‘Higher’ could mean just that: higher, although in research we would mostly think of higher in a statistical sense or even in an economic sense. Stock market returns of 0.1% more than average could mean statistically significantly higher returns in January but those higher returns would economically not be of any importance.

This research question would pass the test as all concepts are well-defined and we do not need to make any other assumptions.

Compare this with another example:

“Do international political crises have an impact on financial markets?”

Your concept ‘impact’ can mean many things. For instance, you may be talking about a contemporaneous impact: if a crisis occurs, do we observe an immediate effect on financial markets? Alternatively, you may be talking about a lagged response: do crises predict future financial market returns? The concept ‘financial markets’ also lacks precision. There are, for instance, stock markets, bond markets and currency exchanges. All this is easy to solve if you are more precise in your formulation (and precision prevents a lot of confusion).

Where disaster really strikes is in the concept ‘international political crises’. One person’s terrorist is another person’s freedom fighter. If you ask 100 people to define an ‘international political crisis’ this will no doubt result in 100 different definitions. According to Murphy’s Law rest assured that the referee’s definition is different from yours.

Even if you first define an ‘International

However, “Sector rotation in different phases along the Business Cycle” is problematic.[4] There are many ways to define a ‘sector’, many ways to define ‘sector rotation’, many ways to define ‘Business Cycle’ and even more ways to define ‘phases’ of that Business Cycle. Can this research be done? Yes, but be prepared to run a large number of robustness tests. You would have to make sure your results are not a result of all your specific definitions. The more potential ambiguity, the more robustness tests you will have to do. You can hear the referees’ comments ‘Why do you assume the early expansion phase is one third of the length of the total expansion phase? Would your result remain the same if you used a quarter and made the other phases a bit longer?’ And also be prepared that no matter how many you do, referees can always think of one more.

In my experience, assume that basically every assumption or variation in a concept you can think of requires a robustness test in one way or the other.

For a good research question you are looking for concepts that are clear-cut and require as few assumptions as possible. Preferably, let others make assumptions for you and, if you have to, use generally accepted assumptions. There are conventions in the literature that everyone uses. (For instance, ‘official’ business cycles do exist in agencies that date the business cycle). However, always keep in mind that the fact that these assumptions and conventions are (widely) accepted and used by others does not make them necessarily right. Sometimes not willing to accept those premises and going against the flow makes for the most interesting new research.

1.1.2 Less is more

The Germans have a saying by Goethe: “In der Beschränkung zeigt sich erst der Meister”. ‘Less is more’ comes close but it basically says that true masters of their craft know how to restrain themselves. KISS (Keep it simple, stupid) may come closer.

Let’s consider the “Are stock market returns higher in January?: an international study” research question more closely. Compare this with “A comparison of five theories of asset pricing applied to the Dutch stock market”.

Both research questions have pretty clear concepts but which one is better from a budding researcher’s perspective? To thoroughly understand five different theories and apply them requires an awful lot of reading and understanding. If these theories are important, other researchers will have written a lot on each of them and to grasp it all will require a lot of knowledge and even experience. Some theories will build on others; some will be linked. Testing these theories may also require a lot of different tests. The next question is whether it is wise to look at the Dutch market. There are around 17 million Dutch and not all have an interest in Asset Pricing. So you cater to a very small audience.

Now look at the January research question again. This is a very easy test (you could do a difference in means test or simply run a regression with a January dummy). Once you have figured out how to download data for say the Dutch market, you can, with just a bit more work, download data for many other countries. With every index you download, you

[4] ‘Sector rotation’ is the phenomenon whereby investors switch from one sector to the other in different phases of the economy.

increase the interest in your work. The main investment in time is in figuring out how to download the first data series. Moreover, whether you run a regression for only one country or forty nowadays hardly makes a difference.

I am not saying testing the five different theories may not be good research. Just ask yourself whether you have the knowledge, skills and tools to do it. Why using Dutch data is a bad idea, I will explain later.

1.1.3 Proxies and concepts

Sometimes you can only use ‘proxies’ for your concepts as there are no data available. Again, let’s consider an example.

“Does vacation behaviour of investors affect stock market returns?”

‘Stock market returns’ is a well-defined concept. The concept ‘affect’ will need more precision, but the real problem is with ‘vacation behaviour of investors’. While it is a clearly defined concept, it will be hard to measure. We would need data on when investors take their holidays and for how long. These data are not available. We can come close if we are willing to make assumptions. We could for instance consider inbound and outbound travel data in different countries or the monthly of number of airline passengers or the number of annual leave days in different countries. You get the idea. We may come close with our ‘proxies’ because all these variables are likely to be correlated with actual investor vacation behaviour, However, they remain ‘proxies’ as we need additional assumptions (investors all show similar vacation behaviour) to make it work. If you use proxies the less they align (or correlate) with what you want to measure the more of them you will need.

If you can only find one proxy to test your research question and it has low correlation with your concept, you may want to reconsider your research question to avoid ending up in trouble down the line. Suppose you find an article in the psychological literature suggesting that social people tend to use 'we' more than 'I'. You now want to test whether more social leaders make better managers (better defined in a clear way). You analyse conversations of managers and measure their usage of ‘we’. You find a correlation between your proxy and management quality. You conclude more social managers are better managers. This would be overclaiming. That statement is too strong because the only thing you can really conclude is that people who say 'we' a lot make better managers. Other personal traits may also lead to higher 'we'-usage (which the psychological study might not have tested). So in this case you would definitely need more 'proxies', or a study that shows that no other traits can lead to higher use of the word ‘we’.

Even in top journals researchers often have a tendency to assume results for their ‘proxies’ hold one on one for the ‘concepts’, incorrectly assuming that they are measuring the real thing. For instance, we do not observe ‘investor expectations’ as such, but very often we will use expectations based on some historical data or some sort of model as actual expectations of investors. Or, we use the predictions of analysts as expectations for all analysts, and these analysts are not all investors. These are, at best, proxies for the real thing and assuming that they are the real thing would, for most of these proxies, be silly.

This is an important distinction and confusion occurs frequently so let me give you another example, the so-called ‘equity premium’. This is the return investors expect above the risk free rate (which is proxied by the short term interest rate) to compensate them for risk. We often use long historical stock market returns in excess of the short-term interest rate as an indication, but that does not necessarily equate to the equity premium. Unfortunately, we do not know how high the equity premium is or even whether it exists (and historical premia can vary dramatically depending on whether you go back 30 or 300 years as many pension funds are currently discovering). Top senior researchers might get away with mixing proxies and concepts, but if I were you, I would not bet on it.

In my view, if you use proxies for concepts, carefully spell them out and where possible refer to your proxies and do not extrapolate this to your concepts. Otherwise, people may feel you are overclaiming. What you find for your proxies may not necessarily hold for concepts.[5] Some proxies are accepted as everyone in your field realizes that it is the best you can do but be aware of their shortcomings throughout your writing.

1.1.4 Make things as simple as possible…

Some academics start out with the misconception that because we often deal with complicated issues, everything we do has to be or at least look complicated. Alternatively, you may be of the opinion that in order to look brilliant you have to make things look complicated. I prefer Einstein’s view: ‘make things as simple as possible but not simpler than that. Brilliance is in making things look simple. This is difficult enough as it is. The biggest compliment I can get is “that I can make things look so simple”. I work very hard at that’.[6]

In my view a good test of your research question is whether you can phrase it in a simple understandable question. Often, if you cannot, you have not thought enough about your research topic.

The adage of real estate agents is ‘location, location, location’. Similarly, academics talk of ‘contribution, contribution, contribution’ or its closely related cousin ‘motivation’. I will use them interchangeably.[7]

---

[5] Many economists in my experience have a similar tendency. They built models based on assumptions that are clearly violated in the real world (and they are happy to admit so) but subsequently they show how their model explains a phenomenon by their model and ignore the possibility that many aspects of the real world they happily assumed away might explain the same thing.

[6] Consider Steve Jobs’ quote in this respect: “That’s been one of my mantras — focus and simplicity. Simple can be harder than complex: You have to work hard to get your thinking clean to make it simple. But it’s worth it in the end because once you get there, you can move mountains.” BusinessWeek, May 25, 1998.

[7] More precise we could define them as follows: **Motivation:** why is what you do interesting and to whom? “Who cares?” to put it bluntly. **Contribution:** What do you add to the literature that we do not already know? These are closely related because if what you want to do has already been done in the literature it is hard to answer the ‘why is this interesting’ question

The first thing editors and referees wonder about when they see your work is why your research is important, and what your research adds to the existing literature. The larger the contribution, the more likely you are to get your research published in top journals. If your research is sound and your paper is well written, than the only thing that determines success is your contribution as seen by the editor and the referees (although a bit of luck might play a role as well). Assuming you can research and write, than the contribution of your work and how to get your contribution across is the only thing to worry about.

Therefore I will discuss contribution at length and although I may bore you to death now, you can send me a thank you note later in your career. Let’s start with what many researchers find difficult - figuring out your contribution.

1.2.1 What is your contribution?

So how to assess the contribution of your research question? A first test is the blunt: ‘Who cares?’ Once your research question is answered, how will your results change lives? And whose lives? Practitioners? Regulators? Academics? How is your result going to change behaviour, views or theory? The more detailed and specific you can make this the better. If you can think of ten different ways of motivating your paper, do it. There may be more than one motivation or contribution, so make sure you get them all. But also make sure you do not overclaim. This often happens if you start mixing up concepts and your proxies.

The ‘who cares’ question is often hard to answer, particularly if you are in the midst of your research. You know it is there, and you can feel your contribution, but it is hard to spell out. Often it helps to try and distance yourself from your research, get away from your computer, walk on the beach or use other means necessary to relax.

Assessing your contribution is especially difficult if you are new to this game but it might always remain an issue. Here are some tricks I have found useful.

‘Who cares?’ if you find the results that you find, ‘Why?’ and ‘What does it mean (it=your results)?’ These are questions you cannot ask yourself often enough from as many different perspectives as possible. So keep asking and keep changing perspective, imagine you are reading your results as an investor, as a regulator, as a fellow academic, as a penguin (if you are working on penguin research) or, if you dare, as a referee.

Revisiting your motivation and contribution every couple of months is another good trick. Your perspective changes over time (and in most cases years can do wonders for the views you hold). As an example, in some recent work my co-author and I looked at returns in ALL stock markets in the worlds using ALL historical market indices (including dividends) and ALL historical short-term interest rates available. We were so engulfed in our research question that we missed that as a by product we derived – what we believe is – one of the best historical upper bound estimates of the ‘equity premium’

---

and if you find it hard to justify what you are doing you probably do not make much of a contribution.

available (if you are in finance : it's 3.7% ☺). That byproduct might for many researchers be more interesting than our main question.

Another trick to tease out your motivation/contribution is to 1) take your research question 2) imagine the outcome you would like to see from your tests and 3) imagine you are a journalist writing a newspaper article about your research. What would your results look like, and what would the article look like? If need be it can be very helpful to actually try and write such an article yourself. Again, a change of perspective (here from researcher to journalist) may help to sort out your contribution and motivation.

How does your research question (and/or results) relate to the literature; what is similar and what is different from what others have done? Do you use different concepts, proxies, assumptions, data, techniques and why? This positioning of your paper is extremely important and I will come back to this later, but writing down similarities and differences is a good start to get a grip on your contribution.

How big does your contribution need to be? When it comes to contribution bigger is better. As noted before, if you can do your job as a researchers, the size of your contribution and luck are the only determining factors as to whether you get a top publication or whether your work is more likely to end up in a lower level journal.[8] So how to assess the size of your contribution?

First and foremost, keep in mind that the problem with contribution is that it is highly subjective. You will learn this lesson the hard way once you get your first rejection. But while referees and editors can be wrong, keep in mind that once you have a research question, its contribution is easily overestimated, particularly by the person who just came up with the question and especially if you are just starting out in this business. Most research questions tend to look really, really good immediately after the inspiration has just hit you.

Here are some reality checks on the actual size of your contribution:

1.3.1. The textbook question

---

[8] This does not mean you should only do top level research. Generally it takes ten times more work to get published in an A* journal versus an A journal (unfortunately most universities fail to appreciate this). So even papers with a lesser contribution - if they do not take up too much time - can be beneficial (especially if you are at a university where they do not realize that A*=10A). Moreover, while you should not shy away from top publications at a younger age, they may be less likely. There are only so many spots available in these top journals. A bit of experience does help and top notch research questions are hard to find (more about how and where to find them later). So as long as you do not need to spend much time on them, even papers that do not have a top contribution might be of interest. There is always a trade-off between time and publication level. As a researcher, time tends to be your most valuable asset (and unfortunately too many people and institutions have no problem frivolously wasting it and that trend seems on the rise). Just to be sure, it may be good to repeat that one: as in life, as a researcher time is your most valuable asset. Spend it wisely.

A helpful tool to assess your contribution is to think of the ‘textbook question’. Look at the main textbook in your field. Suppose your research pans out as planned. Where and how in the textbook would it be discussed in say ten years time? Would it be referenced? Footnoted? If the answer is ‘yes’, you are probably talking top research already, particularly if you are thinking of an undergraduate textbook. Will your paper get a sentence or a paragraph? If you think it might be a chapter start thinking Nobel prize. But if you can tell now that none of the above will happen your paper may not be top publication material.

1.3.2 The relation with the foundations in your field

A second tool (which may also be helpful in coming up with research ideas) requires a bit more experience. But you can give it a go. Sit down and think about the major challenges in your field of research and see how your research relates. Let me give some examples in my field.

One of the fundamental concepts in finance is that risk and expected return should be related. Who would be willing to take on risk if there was no reward? Now, the problem is that while in theory this should be the case, empirically we are not so sure. The problem is that we use all sorts of proxies for risk (like volatility) and also for expected returns (like historical returns). In our field we believe so strongly that this relationship should exist that precious few people realise that the empirical evidence we have may be questionable.[9] Now think about this again in terms of the concepts ‘risk’ and ‘expected returns’ for which we use all sort of proxies like volatility and historical returns. Suppose you can come up with a better proxy for one of these and then show that risk and return are strongly related…. So it may be worthwhile to see whether you can come up with one of the two.

To find your contribution, ask yourself in how far does your research question help our understanding of some of the fundamental hypotheses and relations in our field. Can your proxies overcome some of the main problems with generally used proxies?

Now let’s go back to the research question:

“Are stock market returns higher in January?”

Is this an interesting research question in terms of contribution? Well it used to be. To see why, you need to know another important hypothesis in finance. In the 1970s we thought financial markets were essentially unpredictable based on the so-called Efficient Market Hypothesis (information is efficiently priced in by market participants; only new information moves prices but new information is unpredictable by definition, otherwise it would not be new; therefore price changes are random). The hypothesis was derived from the empirical finding that markets were difficult to predict (yes, empirical results preceded theory). Several studies confirmed this. But if all (publicly available) information was priced in then surely this month’s information should not predict whether returns are higher or lower than they would be in in any other months. So if I find that stock market returns are significantly higher in January, that would suggest some predictability, or in other words, a violation of

[9] See for instance: this video on “Asset Pricing explained”, by Eric Falkenstein. https://www.youtube.com/watch?v=0ugUZzUL0WY

the hypothesis, especially if I show that this is the case in many countries and different time periods.

In a world where people believe that markets are informationally efficient this might be an interesting 'seasonal anomaly'. If you were the first to test it, it would make a contribution. Although it is questionable whether it is still present and even whether it ever really was[10], the January effect is still one of the best known 'seasonal anomalies' to date (Google Scholar reports over 5000 hits).

The January effect made a contribution not because it was interesting per se to know whether January returns were higher but because of the link with the theory. The more you can link results to the main theory and the more explicitly so, the larger your contribution tends to be.

1.3.3 Top papers
What papers are you referring to? If the main papers your work relates to are in top journals, your paper is more likely to have a bigger contribution. Or, maybe I should state this the other way around, if the main article your work relates to is not in a highly ranked journal your work is more likely to make less of a contribution.

1.3.4. The debate
An often neglected contribution to the literature is if your paper adds somehow to a debate. Surprisingly, you often see studies that only focus on discussing the studies in the literature that are in line with the conclusion of the authors. This is a bad idea for three reasons. Science is hardly ever one sided - it suggests that you are biased. The referee will most likely know the literature and be surprised if you only pay attention to studies confirming your results. But most importantly, if there is debate going on (and when isn't there?) adding to a debate can be a huge contribution. Just confirming one side of the debate is probably less so.

1.4.1. Don't do research on Dutch data
Probably the worst research questions are the ones that result in papers that have something like the "New Zealand evidence", "the case of the UK" or the "Dutch experience" in the title. Unfortunately, if your research confirms a well-known phenomenon in your country, it is a so-called replication study and unlikely to get you published in a top journal.

Here is the important lesson. Using data from a country just because you happen to live in or originate from that country is a very bad research idea. Let me say that again. Using data from a country just because you originate from or live in that country is a very bad research idea. Use the conventional data from conventional countries in your field unless

[10] Nowadays, there are many more tests a so-called anomaly needs to pass and it is questionable whether the January effect would have done so. For instance, as there was at the time no reason to believe that January should have been different it might have been a case of what we would now call datamining. In any given sample some months will have higher returns than others and if you have 12 months….

your dataset can tell something about a research topic no other conventional dataset can.[11]

In most finance research, data from the United States are the default data sets. Not because we necessarily like the United States but these are simply the data everyone knows and the three top journals reside there. You could argue that by just looking at the United States we may have all sorts of biases. The US has been a superpower with hardly any military struggle within its borders so de facto we are looking at what might be an outlier. And yes, I would agree with you. You could even argue that the automatic assumption that it is ok to use US data might even make for bad science. Again, I would agree. Unfortunately, most top researchers don't and data from the US are the go-to-data (and generally the best cleaned data around when it comes to other biases).[12]

If you need an alternative, I like to use international data as much as I can (makes the study interesting to as many people as possible). But looking at many different countries is not always easy. Good data may not be always available or may be computationally burdensome if you want to look at individual stocks.

Sometimes some countries have data that allow for studies the conventional data do not. Finland is a good example. In finance there are a lot of studies using Finnish data, not because the Finnish are special when it comes to investing or because we particularly care about Finnish investors or the Finnish evidence but because we can measure things in Finland we cannot measure in the US.

Here is the important lesson: your research question should dictate the data to use and if your research question can be done using conventional data (or a conventional data gathering approach) in your field, use these. If you cannot, then use data that can best answer your research question.

For instance, I have a paper using UK data, but only because we needed a very long data series and the UK data went back to 1692 when the US did not have a stock market. In another paper we use data on New Zealand investors simply because we can measure things about their behaviour we cannot measure elsewhere.

My advice would be that if you now are working on a research idea and are thinking of using something like the 'Malaysian Evidence" or the "Japanese experience" consider ripping it up. You limit yourself to B or C journals if you are lucky. Using the wrong data has been my biggest time waster. I have had some good papers go down the drain because in my very early days some of them were replication studies and later - even though they did something new – because they used the wrong Dutch data and referees felt unfamiliar with them.

[11] Of course, if in your field it is perfectly normal to use data from the country you live in, then it is perfectly fine. However, if you feel a need to add something like the Dutch evidence in the title of your paper, that may be a signal you are doing a replication study and you should worry.

[12] What might become more and more accepted in finance is Chinese data. But if you are a young researcher and even if you are Chinese - or maybe especially if you are Chinese - I would be very careful.

As this is the most common mistake let me elaborate a bit more.

The motivation for a study on the Mauritius stock market I encountered once was ‘We are the first study to look at the evidence from Mauritius’. The fact they were the first to study the Mauritius stock market did not come as a surprise to me. Many people would not know Mauritius existed, let alone that it had a stock market. Unfortunately, that does not make the study interesting. The study tested a specific theory on stock markets. This theory originated in the US where millions of stocks are traded every nanosecond. A more interesting motivation might have been how this theory would hold up in a very illiquid market where a stock hardly ever changed hand (as I would expect could happen in Mauritius).

Sometimes students realize they have the wrong dataset and try to justify it ex post by claiming that they look at how a theory holds up in ‘a emerging market’. Beware that the problem then becomes why they did choose that specific country as an emerging market. There are many emerging markets around so you will need a better, or more specific, reason. So before you pick a dataset think carefully. If Mauritius has a competitor as the most illiquid stock market in the world, people will wonder why you did not use that one.

The country argument also goes for groups of countries or regions. If your research does not specify that a group of countries should be prone to a phenomenon then there may be no point in studying this group. While studying emerging markets might make sense, studying African countries may be less so. Just to prevent confusion, there is nothing wrong with African countries. It all depends whether African countries are relevant for your research question.

Here are some arguments you frequently hear when you ask people why they chose the wrong data.

But I am the first to use Dutch data
Being the first can be good news or bad news.

But I am a Masters student
There is no law preventing Master students from publishing in top journals. Why would you disregard that option from the start? I have been lucky enough to have top publications with some of my Master students all using international data and this has helped them in their careers in the real world.

But not only do I use different data, I also do different tests….
Researchers sometimes argue that the dataset is fine because they are going to add more hypotheses and tests than there were in the original study. Unfortunately, that only makes the water murkier. The reader will be left wondering whether the results are caused by the difference in datasets or by the change in tests? If you have new tests to add, it is better to use the conventional dataset. If you want to do both, then first show what the impact is of the data and then what the impact is of the new tests, method or hypotheses you want to add. (Start with the reference point, test and dataset the reader is accustomed to, if you can. If on top of that you can replicate more or less existing results in your own paper, that shows you know what you are doing and builds confidence that your own results are correct.)

You can introduce more than one change in a paper (and changing proxies/assumptions can lead to interesting follow up studies). However, just like in a laboratory experiment, if you feel that a change to previous results in the literature can have an impact, change only one thing at the time so you can trace what is causing the difference.

But the Dutch are a great people….
"I know, I know.", to quote Graham Norton. I could not agree more. And the world has an awful lot to thank the Dutch for.[13] And more research should be done about their greatness and their influence on the world and the world would be a better place if we all were Dutch. But that still does not mean you should use Dutch data.

But I work in a Dutch university - isn't it my duty to work on Dutch data?
I am all for giving back to society but the biggest gift you can give is top notch research. There are two possibilities, which are both good news. 1) The Dutch are not likely to behave any differently with respect to your research topic. This is good news. You can just take international evidence and generalize that to the Dutch case. 2) The Dutch are special with respect to your research question. Basically good news again. You can use Dutch data and they may also add something to our generic understanding. However, I would be very careful before you do this as a young researcher.

Now there is nothing wrong with doing a bit of analysis on a Dutch dataset for local purposes. But see it for what it is: a bit of consultancy. Do not mix it up with your goal of top research.

1.4.2. Pseudo motivation
Many researchers only pay lip service to contribution/motivation. Let us go back to the vacation behaviour. A paper may start with 'Vacations become increasingly important in society. We investigate the impact of vacations of investors on stock market behaviour'. The problem here is that this does not explain why it is important to look at the link between vacations of investors and their impact on stock market behaviour. You motivate one of your concepts in the research question but not the research question itself: "pseudo motivation". You see this very, very often. Many papers start with the assertion that A, B, C or D is becoming more and more important so this warrants the study. That is fine – although not very original - if A, B, C or D directly goes to the heart of your research question, but not if A, B, C or D goes back to one of the concepts in your research question. To drive this point home consider "Are men more optimistic than women" as a research question and the nonsense of the paper starting 'The population of men and women is growing rapidly and they are becoming increasingly important; therefore we study whether men are more optimistic than women'.

1.4.3. Lazy motivating
Next up is another mistake of lazy motivating as in 'studying the impact of vacation behaviour on investment decisions is important for risk management of investor portfolios'. This is obviously true but the problem is that it is very hard to come up with a research question in finance which is not somehow important for "risk management". Finance is to a

[13] See Russell Shorto's: Amsterdam: A History of the World's Most Liberal City (2013) for an overview.

large extent risk management. So it is a hollow phrase. If you can show how (in detail) vacations affect stock market behaviour and how that translates (in detail) in a way that would affect risk to an investor portfolio, and how investors could manage that risk to prevent this behaviour having a major effect on the portfolio, you are getting the point. The simpler and the more specific the better.

1.4.4. The paradigm shift…..

This is not really a mistake but it may be good to discuss it. Taking my words on the size of the contribution here to the extreme you could think that if you can come up with an earth shaking contribution that would turn current knowledge upside down (and all tests are correct and your paper is well written) you will get published in a top journal.

Beware. Keep in mind that no matter how good your paper is, life in academia is not always fair. Many good papers do not make it into top journals simply because of bad luck. The editor does not like it, or referees do not like it, for no clear reason. Unfortunately it does happen quite often. Research is also sensitive to trends. Some topics are hot and your research may just be on a not so hot topic no matter your contribution. On top of that academics are (almost) human. They have - sometimes incorrect - kneejerk responses. They are risk averse and tend to be sceptical of research that is too confrontational. So ideally your results should conform to some extent to mainstream research but not completely. But if it is completely orthogonal to mainstream be aware that old paradigms die hard.

A good example is the so-called behavioural finance literature (the notion that investors might not be strictly rational, even collectively). No matter how solid and correct the research results might have been in the early days, behavioural finance was considered a paradigm shift and much frowned upon. Nowadays this may seem hard to believe but the early behavioural finance people had to carefully motivate their results from a main stream perspective and carefully align it suggesting that some people might not be rational all the time.

So your research must be different but not too different. The younger you are the less different it should be. But finding the right angle between mainstream and new work can be difficult (I still struggle but maybe because I am too much of a Frank Zappa fan: without deviation from the norm, progress is not possible). This may be frustrating because it suggests that only more mature researchers are allowed to take a 'walk on the wild side' to quote another great musician. But then again, when you think about it a bit longer, maybe that is for good reason.

Having said that, the older I get the more I would argue that one of the reasons we have not made much progress in our understanding of financial decision making in our field is to a large extent based on too much trust in mathematics and a reluctance to accept research that goes off the beaten track. But that is another story.

1.4.5 Doing bad research for the wrong reasons

When you ask researchers about a bad research question some of the arguments that they give are 'at least it is a good learning experience' and 'we wanted to do a paper together'. Life is short and if you must waste it on bad research that is of course your choice. But why not try to do something that is good research and has a bigger contribution instead.

I like to question everything - preferably conventional wisdom - and tend to wonder 'Why?' a lot. While it does not always make for easy living (for me or the people around me) it does generate a lot of research ideas. Of course, ninety percent of these ideas do not make it past the drawing board, but the ones that do have made for some nice research.

So my first advice would be to no longer go with the flow and start questioning everything. I am biased but strongly believe that a nice side effect would no doubt be that it leads to a better world if everyone did so. But not everyone might agree.

If you feel that this is a step too far, start questioning the main principles that currently guide your field. When I started out in investments, I simply could not believe that markets were almost completely unpredictable (the current thinking at the time). In those days fellow academics in finance would raise an eyebrow if you wanted to do research to the contrary. I got quite a few papers out of my willingness to question that hypothesis.

Keep in mind that in academia fashion plays a role; trends come and go and nothing remains always true. For instance, it is generally accepted at the moment that smoking is bad for your health. That may be true but would it be bad for all people all the time? Or could there be examples where the health benefits outweigh the health costs to an individual - if only for some people? How do we measure health as a society? Life expectancy? And is that a good proxy for happiness? What behaviour do people show after they quit and is that really healthier than smoking?

Or, to put it more bluntly in the words of Joe Jackson:

"It is has become 'common knowledge' that smoking is one of the worst things you can possibly do to yourself; 'all the experts agree'. Of course, 'all the experts' once agreed that masturbation caused blindness, that homosexuality was a disease, and that marijuana turned people into homicidal maniacs." (Smoke, Lies and the Nanny State, Joe Jackson, www.joejackson.com)

What may work well is to go back to your textbook and see which topics intrigue you. A similar approach would be to have a look at the top journals in your field and see what issues are of interest. There are trends in research. The downside is that this may be dated (and the trend may have passed). Still, if you can find up to date sources of academic working papers (like SSRN (www.ssrn.com) for the social sciences) that may be a good start. A second problem is that there may be no longer any low hanging fruit in that part of the orchard. If many academics are working in an area all the easy pickings are probably gone. A third problem is that unbeknown to you someone may be working on a similar idea. In most cases in my experience the chances of other researchers coming up with exactly the same idea as you are remote and if they do the approach to that research question is likely to be very different from yours. However, some ideas do have a first mover advantage so you are doomed if you come second. I know that a lot of people use the literature as a starting point but I personally try to look for places in the orchard where no others are searching.

If you do want to pursue an idea already getting traction in the literature make sure you look at working papers or papers presented at conferences and seminars rather than the already published articles.

Daily life, newspapers, Ted talks and practitioner publications may be other interesting sources of inspiration. But with news events you have to watch out. The 2008 Financial crisis stirred a flutter of academic activity in finance. The problem is that events like these also attract the famous and the really smart researchers so it will be very hard to offer a fresh perspective.

Also, while 'case' studies can be interesting, case studies might be perceived by some as a bit 'one observation' research. In my view good research should generalize widely, so the more observations you can rely on the better. With case studies there is always the risk that the results are very specific. I like research with many observations (although we now seem to enter an area where we have so many observations that significant levels seem to become meaningless). But sometimes it can be interesting to do a case study on a specific phenomenon like a company or a crash or a crisis, as long as you keep in mind that you do have only one observation. Always consider if you do a study like this, whether the lessons learned can be generalized.

In many fields new data opportunities manifest themselves on an ongoing basis. Sometimes just having a look at what type of data is on offer, even going back to well known databases, might generate a new research idea. But in the 'deep web' (http://en.wikipedia.org/wiki/Deep_Web) you may be able to find all sorts of interesting data. Sometimes great data maybe freely available if you look for it. Even data on international political crises ☺.

Try to predict where your field of research is heading. There is now more and more research that suggests that financial market returns are to some extent predictable. So just finding a new predictor in itself becomes less and less interesting. Instead the focus has moved to the economics: what is driving the predictability? Alternatively, you need predictability with a twist (for instance time varying), gradually moving to the 'how to predict' and problems related to that (statistical biases, model uncertainty) or to papers that focus on what happens to variables that seemed to predict in the past. At the time of writing, papers in these areas are already popping up.

Try to think of the big problems in your research area and try to make a contribution: 'how to measure risk in other ways than using volatility?' 'Can we measure expectations of investors?' are two questions I ask myself at the moment.

Maybe you can look at even bigger questions. There are a lot of problems in this world. Sometimes – at least in my field – I get the impression that we in academia focus very much on moving the deck chairs and not often enough on how we can steer the cruise liner we are on. Maybe I have been watching too many Ted talks recently.

Be creative.[14] What do you want to measure and how can you do it? Some researchers have gathered garbage data to measure consumption levels. Others have used noise levels in the trading pits of stock markets to measure investor sentiment. And the first researchers getting access to trading results of individual investors are now famous. The advantage is that you might be for a long time the only one working in your field without any competition. In my view there is little creativity around at least in finance so when they do pop up they can be remarkably refreshing even if they use garbage data.

Don't be afraid to take the road less travelled. In today's academic world you hear more and more that our goal is 'a publication'. That is of course incorrect. Our goal is to advance knowledge and publications are hopefully just a reflection of what research is all about: the setting off into the unknown, discovering new things. Unfortunately, today's academic world is so focused on 'publications' that we sometimes lose track of the real thing and the real joy of doing science.

When you think of a research question most of the time you will not know the answer upfront. In science we have a tendency to focus on research that produces statistically significant outcomes. If results are not significant your hard work might have been for nothing. In most cases not finding a significant result means no publication. The best research questions, however, are those where results are interesting regardless of the outcome. A significant result may be a pleasant surprise but an insignificant result would be good too. For instance, suppose that a study some 25 years ago found that some economic variable predicts stock returns. You re-do the exact same study and find results are still significant. That is interesting because results have not been arbitraged away by investors (as economic theory predicts). If on the other hand you find that results are no longer significant then your results show that economic theory did work.

It is hard to come up with research questions where the results are equally interesting regardless of whether you find a significant result or an insignificant one. But sometimes

---

[14] To quote Steve Jobs once more. 'Creativity is just connecting things. When you ask creative people how they did something, they feel a little guilty because they didn't really do it, they just saw something. It seemed obvious to them after a while. That's because they were able to connect experiences they've had and synthesize new things. And the reason they were able to do that was that they've had more experiences or they have thought more about their experiences than other people. Unfortunately, that's too rare a commodity. A lot of people in our industry haven't had very diverse experiences. So they don't have enough dots to connect, and they end up with very linear solutions without a broad perspective on the problem. The broader one's understanding of the human experience, the better design we will have.' Wired, February 1996

not finding an expected result can actually be a good thing, Take your time to consider 'What does this result mean?' and think about it a bit longer. As Isaac Asimov put it: 'The most exciting phrase to hear in science, the one that heralds new discoveries, is not "Eureka!" ("I found it!"), but rather "hmm....that is funny....'

The first thing I ask myself is whether I feel passionate about it. As Hegel puts it. 'Nothing great in the world has ever been accomplished without passion'. If it is a good research question you are going to spend quite a bit of time on it. So is this what you want? What you really, really want? If so….

Start writing your abstract and introduction. More on the 'how to' of writing later but for now write down 1) what you want to do, 2) why it is important and 3) how it relates to the literature. If you need results but do not have them, make them up for the time being.

Then, if you can afford the time put it away and think about it some more (we hardly teach thinking at universities anymore but here is your chance!). Many good discoveries are done when people distance themselves from their topic. Time will put some distance between you and your idea and change your perspective. Rest assured, a good idea will remain a good idea and will remain a good idea no matter from what side you approach it. Ask yourself whether it will get you in a top journal. Why? What makes your idea good enough? And how much time will you need to spend on it? (Generally your research will take twice as long as you think even if you take into account that it will take twice as long as

Now play devil's advocate and start telling yourself why your research idea is a bad idea. What arguments go against your idea? Write them down. Can you tackle them? Imagine you are in court and have to convince the jury that X killed Y. What arguments could your opponent put forward to convince the jury differently? Your results may be consistent with your story but with how many other stories as well? Is what you find a correlation or a (these

academics will always be able to challenge you on something, very often because you are making some implicit assumptions. So down all your arguments and examine them critically and go over every implicit assumption that is in there. Anyone who has done any serious computer programming knows that we make assumptions more often than we think, and very often we only realize how many we have made when things go wrong. Try to tease them out as much as you can beforehand.

Better play devil's advocate yourself before someone else does it.

It is good to think about the implications of different results upfront. Some studies make a good contribution with a range of results, while others require a certain result in order to be interesting. What results render your research of interest?

If you still think this is a good research question go and have a look in the literature again and see whether anyone has done something like this before. What very often happens is that people find out after they have done their research that, somewhere out there, there was this paper….. And while not exactly what they had done, it turns out to be too similar to be able to claim they were the first to do something..... There goes your A*. So after your first search in the literature, go and search again!

Google Scholar is great in this respect because of its search option. But more importantly, as it has the ‘cited by’ option, which allows you to go forward in time and get the latest research close to your topic. This has made the life of many researchers so much easier.

If you are still convinced and passionate find a good (senior) researcher and discuss your idea. Let them shoot holes in it. Ask for their opinion. In my modest view this is what senior staff members in academia are for.

To do proper research you need to be able to gather the right data. If you do not have them yet start searching. Unfortunately, the right data may not always be available and the sooner you know the better. Don’t assume there will be data.

If you can…. Very often a quick and dirty regression and playing around will tell you whether it makes sense to go on. From a researcher’s perspective all alarm bells should go off if you have borderline significant results to begin with. This might become your biggest time waster as it is so tempting to go on. But according to Murphy's Law these results will not hold up. Many borderline significant results are subsequently disproved in follow up studies (or worse once you are in the editorial process working on revise and resubmit versions).

I tend to focus on results I cannot make disappear no matter how hard I try, even when I bias the results against me where I can. Of course the borderline results should worry you even more if they happen in the final results phase.

Your research question should imply the data you want to use. If there is a conventional data set use that. (And if you skipped the worst mistake bit: do not use data from the country you live in or come from).

Do not use ten years of data or five years of data or any round number just because it is a round number. Avoid arbitrary choices. If you do not have to make choice: don’t. If the data go back 11 years, use 11. If the data go back 326 years use 326 years. Arbitrary choices are like assumptions. They require a robustness check. You’ll have enough robustness checks to do as it is. Why add any more?

4.2.1 But the world will have changed in 326 years?

Yes, the world has dramatically changed. The world has dramatically changed since yesterday. So if that is your point why do any research on historical data? The relevant question is whether it has changed with respect to the research question you are addressing. If you are investigating the January effect and you have no reason to assume that this has changed over the last 326 years then use 326 years.

For instance, if you feel that the introduction of a celebration of Christmas changed a January effect why not introduce a before and after Christmas dummy? In the UK people started celebrating Christmas around 1835 (in the US around 1870) so if that should affect a January effect look before and after. If you feel a crisis or another specific period may drastically influence your result, do a robustness test or “dummy it out” (include a dummy in your regression). The same holds for ‘outliers’. Do not remove data if you do not have to. Dummy them out if you feel there is a problem. Or use some sort of robust regression test that makes consistent choices rather than you making the choices. Also keep in mind that in the eyes of some people ‘outliers’ do not exist or, stated differently, everything is an outlier. In fact, sometimes outliers might be the most interesting way to get new research ideas. What causes the outlier? The famous Apple ad is about the outliers and the misfits.

However, it is always good to verify that your results are not always driven by a small number of observations. I find that plotting data and even scatter diagrams to look for correlations are useful tools too often ignored these days. Another useful tactic is to see how your results vary over time. Rather than just looking at two sub periods, verify how robust your results are if you use regressions with a rolling window (and plot your coefficients over time). Nowadays, I see that many students simply do not bother checking some of the basic characteristics of their data or bother plotting some of the fundamental relations they try to measure. As a result I often feel that students are lacking a ‘feel’ for the data, which could have prevented problems further down the line.

Generally, the more observations the better. However, this does not always hold and indeed think very carefully whether your observations are indeed ‘observations’. Let me give you an example (although this may be a bit of a finance issue where we have lots and lots of data at all frequencies). Sometimes researchers feel that if you chop up the data into smaller intervals you create more observations. For instance, you have daily financial data over say twenty years. That will be roughly 5200 (20*260 trading days) data points. This will give you generally pretty accurate estimates. Now suppose you want to do sector rotation over the business cycle. Your twenty years only give you about three business cycles. Unfortunately, this means you basically have three observations. It does not matter if you cut your three pizzas into 5200 pieces. There are still only three pizzas.

To quote Einstein once more: ‘keep it as simple as possible but not simpler than that’. (Figure 1 illustrates one important reason). Many younger researchers like to show off their research skills. They like to estimate using the most complicated stuff possible and preferably put it in the title and use all those big words they came across in their econometric textbook. It’s cool to be a nerd so let’s go for it…. I remember in the good old days when I did a paper on “Temporal Aggregation in GARCH models”. (If you are not familiar with GARCH models or Temporal Aggregation don’t worry but GARCH Stands for

Generalized Autoregressive Conditional Heteroscedasticity, which basically is a complicated way of saying that some time series tend to show turbulent and more tranquil periods). To make matters worse (as I will explain later) we also put it in the title, which assured that absolutely no one was interested in the paper.

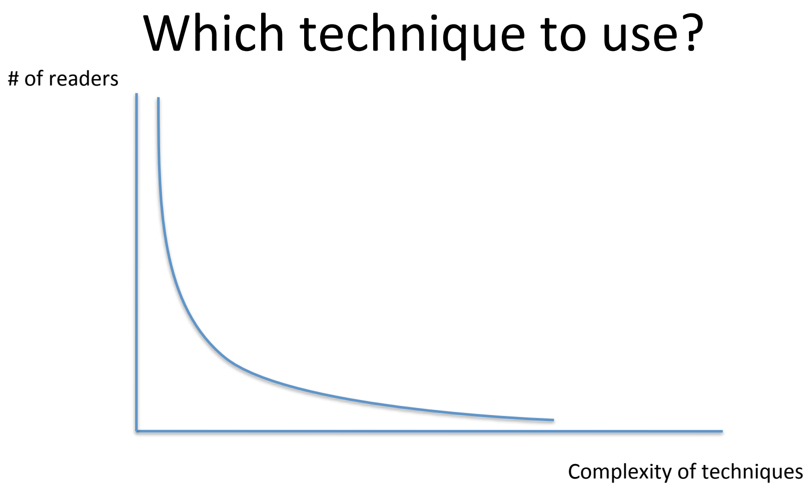

Figure 1. How many readers will be attracted to my papers?

There is a certain appeal to econometrics/statistics and the beauty and sophistication of these models seems to generate a false sense of security and accuracy. However, for most of us (as users) it is good to keep in mind that these are only tools and that our focus should be on what we use them for. There are many different and elegant ways you can put a picture on the wall (but in most cases a hammer and nail will do the trick just fine). In the end, our focus as users of the tools should be on what is in the picture.

Slightly wiser these days I tend to prefer the simplest technique possible. I had dinner a couple of years ago with someone much smarter than I was who understands basically every econometric technique you and I can think of. After years in the field he basically no longer trusted any result that would not show up using a simple regression.

The advantage of simple techniques, like measuring correlations, differences in means tests, or regressions, is that everybody understands them. That means that everyone can follow your story and signal possible flaws in your thinking. Moreover they tend to be very general. OLS regressions (with White standard errors or Newey-West standard errors) are hard to beat and more general than estimating say a specific GARCH model (which may be misspecified if you do not have the correct one).

So if you can use a simple model, start simple. If you want to show that your results are also robust to variations in estimation techniques, use them but put them in a robustness section or in the appendix.

Of course at some stage you may need to use more complicated techniques simply because the simple techniques cannot answer your questions. For instance, if you are interested in the joint effect of both a January effect on the mean and the volatility you will need to estimate those effects with a GARCH model at some stage. But even then I would first consider the separate effects so you can take the reader by the hand and introduce each effect one step at the time.

The next question may seem obvious and obsolete but I am going to ask it anyhow: are you testing what you think you are testing? You'd be surprised how often this does not happen (and I have to admit is has happened to me…..). A research question that asks whether A and B are different is a good illustration of where such problems occur: in the tests the researcher shows that A is significantly different from zero and also B is significantly different from zero but forgets that the research question was whether A was different from B. In other words you need to test whether difference between A and B is significant.

We had a paper once with the title "Are men more optimistic or women more risk averse?" It sounded good but as the referee correctly pointed out, while we did test the first part we did not test the second. Oops, it does not boost your confidence if that happens.

So once you have written down your test, it may pay to double check. The best approach to getting your test right is to first to write down what you want to test in as plain English as possible. Then, link it directly to the equation you intend to use to run your test, trying to make that link as tight as possible.

Let me illustrate this with the January effect - which is easy - but you will get idea.

What you want to test is whether there is a January effect. Formulated more precisely: are average January returns significantly higher that average returns in all other months in your sample?

The difference in means test is obvious but let us look at the regression with a dummy variable for the month of January:

$$r_t \qquad + \alpha Jan \qquad \varepsilon_t$$

where r_t is the stock market return in month *t;* μ is a constant Jan_t is 0 if a month at time *t* is not January and 1 if the month is January, and ε_t is the error term in month *t*.

The error term has mean zero. Now if your dummy is zero, returns are equal to μ. μ is the average return in the non January months (or all other months). If your dummy is 1, average returns are $\mu +$. Therefore $\mu + \alpha$. is the average January return in your sample. We are interested in the difference between average January returns ($\mu + \alpha$) and the average returns in other months (μ) which is therefore equal to α ($\mu + \alpha -$). If you find an α of 0.02 or two percent, January returns in your sample have been two percent higher than average return in the other months of your sample. So the interesting question is whether α is significantly different from zero. If it is, there is a January effect.

This explicit linking exercise may seem over the top, particularly in this simple case. But it makes it very clear that you are doing the right thing. And if you do think this is over the top, keep in mind your reader might find it very helpful.

Note also, that it gives you a feel for the order size of your results. If you find α equals 20% it is very likely you must have made a mistake somehow.

We once got a referee report where the referee had gone wild. The referee basically required us to include every control variable known to mankind (and when it comes to finance, mankind has come up with quite a few). That was silly but they did list all control variables and why not use them all: better be safe than sorry?[15]

The lesson is simple: control for effects that might correlate with the effect you are trying to measure. If there is no reason why your effect (proxy) should correlate with an effect (proxy) you really need not to control for it.

A lot of researchers seem to have difficulty deciding which control variables to include. Papers that have the wrong control variables and lack the proper control variables can be found in all the top journals (don't get me going). A favoured option seems to be to include the ones that other researchers have included. That might protect you somehow against silly kneejerk responses from referees, but it might not make for good science and could annoy better referees. So spell it out if you include variables for this reason. It shows that at least you know what you are doing and why you are doing it.

Include control variables to control for already known alternative explanations which may cause your effect. For instance, if you know that there is a January effect where stock market returns are higher than average in January and your hypothesis is that winter returns (let's say November through April) are higher than the remainder of the year, you will need to control for the January effect.

Beware of the fact that some controls will be elephants in the room. You must control for them ASAP in your analysis and tell the reader clearly that you are going to this ASAP in the introduction. This to show the reader that you know what you are doing. Because until you do, they will keep wondering. For instance, if you believe that winter stock market returns are higher than summer returns you want as soon as possible to rule out that this may be caused by the January effect. You also want the reader to know that you know! Editors get papers from good and bad researchers so they try to sort out quickly which category you are in.

I generally like it when an effect shows up without any control variables. This makes

If your effect shows up only after you control for some other effects, start there. But then very carefully explain the link between the controls and the effect you intend to measure. Also explain why it might not show up in an exercise where you would not control for these other effects.

If you have followed the guidelines above you probably have noticed it almost suggests a natural way to do your research. It also suggests a structure with your readers already in mind. Start simple in terms of techniques and variables. Try and show your main result/base case as cleanly and simply as possible. Give yourself and your readers a starting point and reference points along the way that they can understand. For instance, if you intend to change some assumptions in a previous study with almost similar data, first replicate the earlier results than add the changes. If you give the reader and yourself a reference point from which to start your journey that makes life for the two of you much easier. Then add complexities (preferably change only one thing at the time) and carefully think and write down what is different to the base case. This ensures you keep your eyes on the ball and later makes it easier to tell your reader a coherent story.

There seems to be weird misconception that numbers are exact. Mostly, they are not of course. Most numbers as we use them in daily life are just statistical concepts or proxies. If you find that the α in your January regression is 0.020134526 and significant with a t-stat of say 2, there is no point in writing it down with great precision. Basically, what you are saying is that January returns are significantly higher and, as an indication, at around the 2 percent a month level. All the other digits are probably meaningless and even the 2 percent may be overdoing it. What you find is that January returns are a bit higher in your sample and may also be so in the future. If you want to get a better feel for the actual value, calculate the actual confidence interval. (You could argue fewer people would be misled if we all reported confidence intervals rather than actual numbers).

Some researchers do not stop short of interpreting a number even if it is not significant (zero is within the confidence interval), particularly if it confirms a favourite hypothesis. If a number is positive or negative but not significant, assume it is what your statistical test says it might be: zero. Don't use it for any inference, it looks silly and the closer to zero your t-stat is, the sillier you look if you start to interpret your number as if it has meaning.

Also be wary of t-stats that are very high. My thesis supervisor used to say that any t-stat over six was wrong - although I have to admit that was before panel regressions became popular. Somehow panel data have done wonders for t-stats. Significant effects galore. But my supervisor had a point. In normal regressions be careful. High t-stats should make you suspicious. The same goes for high $R^{2'}$s.

If your test is ludicrous so are your statistics. I have seen quite a few presentations in my life. Some of them had tests in them that where absolutely ridiculous. However, I have seen many a good academic even after having had it pointed out that the test was bollocks starting to interpret the t-stat of this test as if it still had meaning once the results slide came out. Academics seem to find it very hard to resist an interpretation once they encounter a significant number. Always keep in mind that a t-stat is as meaningful as the tests it relates to. And even if your test is meaningful, a significant result can be spurious.

Somehow, even though they should know better, some academics seem to ignore the possibility that significant results are spurious but immediately cast their eye on the one significant observation out of the twenty numbers or so on the slide.

Very often your results may not only be ‘consistent with’ the explanation you give but also ‘consistent’ with many other explanations you may have not explored in your paper. Some researchers seem to find it hard to accept that possibility once they have found significant results. They tend to become strong believers in just that one explanation. Keep an open mind. It is called ‘re-search’ for a reason.

This concludes my two cents worth on research. May you have fun on your research journey!

Good references (lots of it from the JFE website):

Tips of Rene Stulz:
http://jfe.rochester.edu/tips.htm

Joint Editorial Comments

http://jfe.rochester.edu/jointed.htm

http://jfe.rochester.edu/jointed13.htm

Paper submission check list

http://jfe.rochester.edu/checklist.pdf

Editor of the AER:

http://vita.mcafee.cc/PDF/EditorExperiences.pdf

Writing tips from John Cochrane

http://faculty.chicagobooth.edu/john.cochrane/research/Papers/phd_paper_writing.pdf

Pitching Research Robert Faff

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2462059

From the editors: Publishing in *Academy of Management Journal. Seven Part series. The first one is Topic Choice and appeared in 2011, Vol. 54, No. 3, 432–435.*

http://people.few.eur.nl/kole/WritingAdvice.pdf by Erik Kole

Deirdre McCloskey: Economical Writing, Waveland Press Inc. 2000.

Ben Goldacre: Bad Science (visit http://www.badscience.net)