

Stephen Schaefer

On the origins of factor investing

Stephen Schaefer is Professor of Finance at the London Business School and a co-author of the influential 2009 report on the Norwegian Government Pension Fund and factor investing. In our Great Minds series – a set of interviews with renowned academics and investment experts – we spoke with him about how factor investing came about and went on to become a popular investment approach. We also talked about the major issues researchers should focus on going forward.

Great Minds

You have worked with some of the most prominent academics in the field of finance of the past few decades. Through your work for the Norwegian Oil Fund you have a significant impact on the adoption of the concept of factor investing. How did this concept emerge? What exactly were academics looking for at the time?

“The term factor investing may be relatively new, but the ideas that underpin it have been around in different forms for quite a while now. Firms such as Dimension Fund Advisors (DFA) have been making use of concepts such as the premium on small firms for several decades. Taken individually, this and other findings such as the value premium were initially seen as small steps in the academic field of finance, and we certainly did not imagine they would become so influential on investment practice. Actually, it is only over the last ten years or so, that these concepts have caught on in a substantial way.”

“Let me tell you an anecdote that will illustrate what the situation was like 30 or 40 years ago. One of the very first anomalies to emerge in the literature was the small-firm effect, which was first discovered by Rolf Banz back in the late 1970s. I happened to be visiting the University of Chicago at the time and attended the seminar where Rolf first presented his results.”

“There are two striking things in this story. The first is that, initially, Rolf was not looking for a potential small-size effect in returns at all. He was trying to do something quite different and happened to rank firms in a way that nobody had done before: by market capitalization. This is how he discovered the extraordinary

fact that the returns of small-capitalization stocks in his data were on average much higher than those on large-capitalization stocks and by an amount that was easily as large as the equity market premium.”

“The second remarkable thing, which at the time seemed totally reasonable, was the reaction of the audience. It was a very distinguished group, but their initial response was that Rolf had made a programming error, that he should go away and correct this technical mistake. Of course, there was no error and people finally took his conclusions on board.”

What have been the major changes in the way factor investing is perceived over the past few years?

“What has changed over the past decade is that individual factors are now increasingly considered as part of a broader family. Instead of being enthusiastic about, say, investing in small firms, or investing in value stocks, investors are getting used to the idea of somehow exploiting these investment strategies as a family and looking at the interaction between them. This is actually quite an important element and this is what is actually new.”

“In the early days it was natural to think about anomalies in the context of the Capital Asset Pricing Model (CAPM), in which many people at that time still had a lot of faith. Then, in an important paper, future Nobel prize winner Eugene Fama and co-author Ken French made things much clearer by showing that average returns on a stock were strongly related to its size (market capitalization) and its book-to-market ratio, in other words to its characteristics, and not just to beta.”

Stephen Schaefer

is a Professor of Finance at the London Business School and the author, together with Andrew Ang and William Goetzmann, of the 2009 report, ‘Evaluation of Active Management of the Norwegian Government Pension Fund – Global’ (the ‘Norwegian Oil Fund’), which represented a major breakthrough for factor investing. Throughout his career, he has published extensively on a variety of topics such as fixed income markets, risk management, credit risk and financial regulation. In particular, Stephen Schaefer’s academic work includes a study on corporate debt default in the US over the past 150 years, which was awarded first prize in the 2011 Fama/DFA Award for the Best Paper Published in the Journal of Financial Economics in the areas of capital markets and asset pricing. He was formerly a faculty member at the Graduate School of Business at Stanford University. He has also been a visiting professor at the Universities of British Columbia, California (Berkeley), Cape Town, Chicago and Venice, where he was recently awarded an Honorary Fellowship. Today, he is the lead Academic Director for the AQR Asset Management Institute at LBS. Outside academia, Stephen Schaefer has consulted for a variety of financial institutions. He has also been an independent board member of the Securities and Futures Authority; a senior research advisor to Moody’s KMV; a trustee-director of Smith Breeden Mutual Funds and a member of Moody’s Academic Research and Advisory Committee.

'The initial objective was to say something about the CAPM and expected returns'

Great Minds

"Then, in a slightly later paper, they showed that average stock returns were also related to their sensitivity to portfolios constructed on the basis of size and book-to-market. This was a decisive step from the perspective of asset pricing theory and gave a much richer description of the pattern of return premia across stocks. And this second step was also extremely useful from an investment point of view."


You mean, before it eventually turned into a popular investment approach, all the theory around factors had more to do with uncovering which elements affected expected returns, right?

"Correct. The initial objective was to say something about the CAPM and expected returns. And these findings were extremely important because they showed that the CAPM failed quite badly. It turned out that the value minus growth portfolio generates a big positive premium. So, value stocks achieve better returns than growth stocks but, at the same time, have a lower beta. This is a robust finding and goes against the most important prediction of the CAPM which is that expected returns are positively related to beta."

"From then on, and for a variety of (often practical) reasons, academics started to acknowledge the need to adjust the expected returns they were measuring for their exposure to factors such as size, book-to-market and, more recently, momentum. Otherwise, standard statistical methods in finance such as event studies, simply didn't give the right answer. However, even if this adjustment in expected returns is correct, it doesn't have direct investment implications."

"Take, for example, the 3-factor model which was introduced by Eugene Fama and Kenneth French, back in the early 1990s. As I've mentioned, this model shows that expected returns do not simply depend on the sensitivity to market moves (beta), but also their sensitivity to portfolios chosen on the basis of size and book to market. However, as a pricing model this has no more

¹ C. Harvey, Y. Liu and H. Zhu, '...and the Cross-Section of Expected Returns'



direct implications for investment than the CAPM does. The CAPM just predicts that the higher the beta, the higher the expected return. The portfolio you choose to hold then boils down to how much beta risk you wish to take. And, as a pricing model, the implications of the three-factor model are essentially the same."

"To me, the fork in the road occurred when some people started to regard factor returns as anomalies – inefficiencies that you might want to exploit – while others just saw them as part of a broader asset pricing model. The former is consistent with the way that most investors see factor investing: regarding the factors as potential sources of return that, in principle, everybody should exploit."

"Alternatively, factor premiums could be market compensation for types of risk that the CAPM does not address. That is both perfectly logical and theoretically plausible. A number of alternative models that try to capture some of these additional risks have been proposed. Up to now, however, nobody has been able to translate the factors that have emerged from research on anomalies into the kind of risks that, for example, Robert Merton was talking about when he introduced an extension of the CAPM

that gave rise to multiple factors. In his framework some investors, depending on their own preferences, would be willing to bear these factor risks and others would wish to take the other side of the trade in order to hedge."

How do you see the success of factor investing after your report on the Norwegian pension fund was published in 2009? How has this report changed the way people perceive investment, in your view?

"Well, I am not sure you want to put this in the interview, but let me tell you another anecdote. Approximately five years after the publication of our report on the Norwegian pension fund, a major US investment bank invited me to talk about factor investing at one of their events. I thought: 'Sure!' And then they showed me some kind of booklet they had been using to pitch factor investing to potential clients. Well, on page two of this document was an extensive summary of our work on the Norwegian Oil Fund. So it does appear that, somehow, our work has been genuinely influential, although I did not know exactly to what extent for several years."

"One crucial conclusion we came to was that a significant fraction of the Oil Fund's outperformance, most of it really, was simply due to exposure to the very standard factors we have been talking about. But since these exposures were not built into the benchmark that the fund was using, it was viewed as 'outperformance'. Divergence in factor exposure between the portfolio and the benchmark has important consequences for the way investors should view the performance of their asset managers nowadays."

"Before the crisis and before the large negative returns relative to their benchmark in 2008, I think people viewed the Oil fund as an actively managed fund that was producing small excess returns relative to the benchmark with relatively low risk. Indeed, I think that's the impression you get from reading the Fund's own reports prior to the global financial crisis. The recommendation in our report – which they did not actually follow at the time – was: 'if you want exposure to a given factor, that's something you should decide: it should be in the benchmark.' In other words, these exposures should not happen by accident, they should be the result of explicit decision by the fund sponsor."

"There are two reasons for this. The first is that different factor exposures are accompanied by different types of risk, and the fund sponsor should decide whether or not these risks are acceptable. The second reason is that, if the fund decides to take on such risks,

nowadays there are relatively cheap ways to get different types of factor exposure. Therefore, a fund should only pay an active management fee to someone who can actually outperform a benchmark that actually takes into account exposure to these factors."

"This is crucial because we now know that portfolios with different characteristics – for example a portfolio of value stocks versus growth stocks, or a portfolio of large-capitalization stocks versus small capitalization stocks – will often have quite different factor exposures and will, therefore, often behave quite differently over time. Since we can measure these sensitivities fairly accurately it would be a mistake to use benchmarks that do not reflect these differences."

More specifically, what is your view on the current frenzy of research around factors and smart beta product launches?

"Well, there is obviously a danger about it. We all work on the same data, and the time period we look at increases very slowly, so it is difficult to get new data very quickly. There is a paper¹ you may be familiar with, written by Campbell Harvey, Yan Liu and Heqing Zhu, in which they count the number of factors that academics have identified; there are about 300 of them. They also point out that, because results are typically both identified as 'interesting' and published on the basis of statistical significance, at least some of these results – perhaps many – are likely to be 'false positives'."

"This 'p-hacking' risk is a widespread problem in many areas of scientific research; it is not restricted to finance. Here is an illustration. Let's imagine two completely independent series of returns data, i.e., with a true correlation of zero. If we generate 500 such pairs of data then, even though the true correlation is zero in each case, we would expect to find that, on average, 25 of them (5%) will appear to be significant at a confidence level of 5%. And if, from that finding, I deduce that I have found an interesting and meaningful pattern in 5% of the data, I am making an obvious mistake. In this case the mistake is obvious because I can see the 95% of the data that produces no significant correlation alongside the 5% that appears to be significant."

"So, in this case, I can easily see that it's just luck. But now imagine that 500 different researchers do exactly the same thing with each researcher analyzing just one pair of series. The (on average) 5% of researchers who find a result that is significant at the 5% level are likely to be convinced they're on to something! And, because they appear to have found something, these are the

results that are more likely to get published. We'll never see the 95% of analyses that resulted in nothing at all because analyses that find nothing don't get published. This is simplified version of the point being made by Campbell Harvey and his co-authors."

"You may remember this famous phrase by the 20th century physicist, Sir Arthur Eddington, who is famous for carrying out observations of a solar eclipse in 1919 that confirmed Albert Einstein's predictions about the effect of gravitation on light. He once said: 'it is a good rule not to put overmuch confidence in a theory until it has been confirmed by observation.' But then he went on: 'it is also a good rule not to put overmuch confidence in the observational results that are put forward until they are confirmed by theory.' In other words: if all you have got are some observations and you lack the theoretical framework to understand them, you really need to be cautious."

"But let me make another comment on the current explosion in research and product launches. I think that one of the reasons behind it is linked to some kind of widespread sentiment that future returns may be lower than they have been in the past. And added to this is the fact that you need a huge amount of evidence to be lucky enough to identify an active manager that will outperform in the future and not charge you a fee that will basically wipe out all that outperformance. You'll be fortunate if you find a manager like that."

"There is a growing skepticism among some asset owners, correctly so in my view, about the benefits of active management and about the fees that usually go with it. In a way, factor investing falls somewhere in between conventional active and purely passive management. It differs from passive investing in the sense that it is not purely holding the market. But, at the same time, it is not classic active management, in the sense that the strategies are quite transparent."

"So, for investors who are tempted to give up active management and go for purely passive strategies, factor investing is a very natural thing to look at, because it features many of the elements that can be found in active management."

Are there any other kind of warnings or recommendations you could give to researchers and investors?

"Yes. As I already mentioned, I think we all need to realize that while it is true there is a lot of evidence on the widespread nature of factors, and the historical premiums attached to them, there is not, so far, any fully convincing explanation of why this is the

case. That is not to say that these premiums don't really exist, or that a good explanation isn't possible. It's just that, so far, we haven't got it. In any case, as researchers we should definitely keep looking for an answer to this question. And, as investors, we should always keep in mind the fact that, to this point, we don't have a clear answer to the question."

"In the case of market risk premium it's not at all difficult to explain why there should be a premium on equities as a whole relative to government bonds. Equities bear much of the business risk in the economy and if we could be confident that government bonds would achieve average returns similar those of the equity market we would all hold government bonds. So, the logic for a market risk premium in equities is quite clear. Determining the size of the premium is another issue, but the reason it exists is clear."


"But when it comes to focusing on a specific premium that reflects just one aspect of the equity market, say the premium on value stocks over growth stocks, things become much less obvious. It is not that we can't think of possible reasons for this kind of premium, but there is currently no consensus on what those reasons actually are. And I think investors should definitely be aware of that."

What are the most important aspects or questions academics should focus their research on?

"Let me point out two important areas where I think further research is much needed. As we have already discussed, there is now a very long list of factors in the academic literature; a sort of 'zoo' with a lot of different species in it. I suspect that many these species are related in some way, and if we were able to get a better grasp on the relationship between them, that would help quite a lot."

"For example, momentum and value are generally negatively correlated, and this is in itself very surprising, because both factors apparently have a positive risk premium. Finding an uncontroversial explanation for this phenomenon would be very useful in helping to understand the interrelationships between the different factors."

"The second area of research I would mention is the quest for a genuinely convincing theory of why these factor premiums exist. Until we find some satisfactory explanation for what is going on, I think there will always be a concern that factors may not be permanent and could, therefore, disappear at some point." This is a possible scenario if everybody were to suddenly start investing using factors, right?



'Until we find some satisfactory explanation, there will be a concern that factors may not be permanent'

limit the development of factor investing?

"Up to a point, because collectively we do hold the sum of whatever is out there in terms of assets – let's call it the market. So factor investing is about people deviating from this average position and we can't all deviate from the average position in the same direction. But I do think factor investing will remain successful and will very likely grow. This framework has been around for over 30 years now and it is a fact that the risk characteristics of different categories of stocks are predictably different. So I think that as investors become increasingly conscious of the fact that there are different risk-return characteristics within the overall market, they will take these into account and then make well-informed choices."

"Some investors may consciously decide not to explicitly target any factor premium and simply go for conventional indexation. But other asset owners, who previously allowed their active managers to determine the portfolio's factor exposures, will increasingly make these choices themselves. Some will do so based on their own risk preferences, some will target the premiums attached to the different factors and others will pay attention to both the risk characteristics and the premiums." "So even though there is still an ongoing debate on the actual size of the different premiums, I think factor-based approaches are likely to become an increasingly important part of the investment landscape going forward, just because they enable investors to make more deliberate choices about the kind of risks that they are prepared to take. As we understand more about these risks it will make it easier for asset owners to decide whether they are prepared to be exposed to certain risks and the level of exposure they feel to be acceptable."

"As for the potential limits to the development of factor investing, well, as I have mentioned, most of these strategies imply long-short positions and not all investors can implement the same long-short position. This could limit the expansion of factor investing unless there are enough investors happy to take the other side of these bets. But for now, at least, I think we are still quite a long way from this limit."

"Yes, definitely. There are already concerns about the amount of money that is currently being invested in strategies of this kind, and the consequences this may have on expected returns. And some of these strategies are very puzzling. Think, for example, about momentum. This factor is among the most perplexing, because it leads to an investment strategy that your five-year old nephew might come up with. Why on earth should we invest in things that have gone up and sell things that have come down? I mean, it sounds so naive..."

"So, in this case, it is quite difficult to come up with a rational risk-related explanation. Why has the momentum effect not just disappeared? People have had long enough to arbitrage it away. As a consequence, we cannot simply rule out a risk-related explanation. If a factor is related to fundamental risk, it may well never disappear. And this is why, from my perspective, it is so important to get a better theoretical grasp on what these factors exactly are. And that should be a concern not just for academics but also for practitioners."

Where do you see factor investing going forward? Do you foresee a similar success to that of passive strategies? What would

