Adam Butler:	00:00:02	Hi. Welcome to the ReSolve Institutional Podcast. Today I am really excited to have Robert Carver on as a guest. Some of you may recall that Robert delivered a webinar on portfolio optimization under uncertainty a few months back, which was extremely well-received. Most of you probably know who Robert is, but for those who don't, Robert's an independent systematic futures trader, writer, and research consultant, and is also currently a visiting lecturer at Queen Mary University of London.
Adam Butler:	00:00:40	Formerly, Robert was Head of Fixed Income, and eventually a senior research fellow at Man AHL. He's the author of "Systematic Trading: A Unique New Method for Designing Trading and Investing Systems", and "Smart Portfolios: A Practical Guide to Building and Maintaining Intelligent Investment Portfolios". Robert has a bachelor's degree in economics from the University of Manchester, and a master's degree, also in economics, from Birkbeck College, University of London.
Adam Butler:	00:01:10	Rob, welcome to the show.
Rob Carver:	00:01:12	Hi Adam, it's really good to be back here with you guys again.
Adam Butler:	00:01:16	Yeah, I think this is going to be a lot of fun, and I think we've got lots of really interesting ground to cover. Maybe you can start off just by giving a brief introduction to your career trajectory. I know you spent some time at Man AHL. What did you do before that? What did you do at Man, and what have you done since?
Rob Carver:	00:01:39	I sort of started off my early career not in finance at all, working in the Middle East, which is where I was actually brought up as a child. Then I came out to the UK when I was in my early 20's to go to university, to get a degree in economics. My first job straight out of university was working in investment banking as an exotic fixed-income options trader. I did that for a couple of years, and then I was spending a couple of years doing something different again, which was working in an economics think tank. That time I also got my master's degree. And then I was very lucky to get the job with AHL.
Rob Carver:	00:02:24	My first job with AHL was putting together a new type of strategy for them, which was a systematic global macro, trading across all the asset classes. And then a few years in, they restructured the research and portfolio management teams and put me in charge of the fixed-income portfolio. I ran that until 2013 and then decided to head off on my own, manage my own money, spend more time with my family, all the usual cliches.
Rob Carver:	00:03:00	Since then, I've written a couple of books, which you've already alluded to. I did my own independent research, which I blog about on my website. Obviously - kind of researching trading strategies that I've tried for myself. Recently, the last couple of years, I've also been working as a part-time lecturer.

Adam Butler: 00:03:29 Fantastic. Yeah, you've obviously been keeping really busy. Just a quick audit of your blog, even over the past year or so, you've got over a dozen fairly lengthy articles, some of which we may have a chance to discuss today.

The Focus on Uncertainty

Adam Butler: 00:03:46 I'm going to dive right in, Rob. One of the things that I noticed as I was reviewing your material, and some of your older blog posts and your books, was this deep omnipresent, internal feeling of this focus on uncertainty. It infuses all of your writing. It infuses the presentations that you've posted online.

Adam Butler: 00:04:20 This is something that in my experience, this embedded sense of uncertainty and risk, is rarely approached from an academic perspective, or from a theoretical standpoint. Usually, people internalize this concept of uncertainty by getting a few frying pans to the face, right? Having some really negative experiences in trading, or in some other dimension of your life.

Adam Butler: 00:04:49 Can you point to any experiences, or just things along the way that helped you to develop this heightened sense of uncertainty?

Rob Carver: 00:05:03 I think it's something that probably gradually developed over the years I spent working at AHL. And it kind of came to a ... well, I'm not sure I'd call it a crisis point, or a turning point, but in the last sort of six months before I left, I was running the fixed-income portfolio, and there was a big turnaround in fixed-income prices and yields, which, those of you who were trading at the time might remember, there was discussion about whether the Fed was going to change its quantitative easing policy.

Adam Butler: 00:05:37 This is the taper tantrum?

Rob Carver: 00:05:39 Exactly, yeah. There was a big sell-off in yields. I think, from memory, it was 100 or 150 basis points in 10 years, over a relatively short period of time. Our strategies were quite badly positioned for that. Quite recently, there was a lot of internal soul searching, and head scratching, and also external pressure from our clients to say, well, essentially saying things like, "Well, we think that something is going to happen", where something was usually along the lines of, well, clearly interest rates are going to now rise, this is the end of the zero interest rate period that we've been in for a few years now, post financial crisis. It may even be the end of the massive secular bull running in fixed-income yields that kind of goes back to, I don't know, like the early 1980s.

Rob Carver: 00:06:36 What should we be doing about this? It struck me at the time, it was a very ... these statements are being made on an almost a complete lack of self-awareness about the difficulty of forecasting or predicting. I also find it ironic that we work in an organization where we were basically assuming that computer models and

systematic algorithms and approaches would do a better job of forecasting the future than we could.

Rob Carver: 00:07:09 But we still felt that we were in a position to make what I guess you could call meta forecasts about the actual performance of the trading strategies themselves. People were saying very confidently, "Well, it's obvious that this will happen. And therefore it's obvious that, for example, this particular fixed-income trend following, or carry strategy, will do very badly."

Rob Carver: 00:07:35 I spent some time exploring this and discovered that, if you actually do this properly as a statistical exercise, in simple terms you could do something like say, "Well, let's partition the past into different states of the world when, say, interest rates were low, and interest rates were high. Interest rates were rising, and interest rates were falling." And having done that, let's then see how these strategies do. How does trend following do? How does carry do? How does the short end of the yield curve do? How does the long end of the yield curve do?

Rob Carver: 00:08:13 Obviously, this is quite fixed-income specific, but it's something you can do with any strategy, or any asset class, or any instrument. This idea of looking at conditional returns. All trading strategies are based on conditional returns. Trend following is based on the assumption that if we condition on the returns of say the last six months, or the last 12 months, then we can get some useful information about what might happen next.

Rob Carver: 00:08:41 I did find, indeed, that if I do this exercise, then an environment where interest rates were rising was indeed not great for fixed-income trend following, fixed-income carry strategies, particularly in certain areas of the curve. But I also found out that the uncertainty around that, the central prediction, was really large, and that actually you couldn't really see, in statistical terms, any strong evidence that a step 1 state of the world was particularly bad, versus the other.

Rob Carver: 00:09:17 To boil it down to an actual 'What shall I do about this?' statement, the what shall I do about this was nothing. Because even if you could predict the future, first of all it's obviously really hard to predict the future anyway. That's where we normally let our computers do that for us. But secondly, even if you could predict the future, actually the uncertainty around the performance of different trading strategies in those future environments was so wide that at best you might say, "Well, I could tilt my portfolio slightly one way or the other." But absolutely is there insufficient evidence to, say, completely deallocate from X or Y or Z.

Adam Butler: 00:10:02 Yeah, and I see that you've got ... I remember reading this last year, but you wrote a pretty lengthy and comprehensive article on this topic, it turns out on my birthday, actually last year 21 February, called “CTA Allocations, QE, Meta-prediction, and Conditional Return Distributions”, which I think walks through a fair amount of your thought process and analysis, and probably I would recommend this to readers who want to learn a little bit more about that.

Handcrafting

Adam Butler: 00:10:37 That was an instrumental experience for you in learning about the importance of considering uncertainty in all of our decision making, and our design process, and all that sort of stuff. I'm curious, because another pervasive concept throughout your writing is this idea of issuing more formal optimization methods in favor of something that you call 'handcrafting.' Can you give us a sense of what you mean by handcrafting?

Rob Carver: 00:11:16 Yeah, this essentially is a solution to the kind of classic problem of portfolio allocation, either between underlying assets like stocks in the S&P 500, or indeed between trading strategies. I'm kind of uncomfortable with optimization as a technique, because it suffers from some flaws which are fairly widely known. It's very unstable, so given, for example, quite small differences in the inputs into it, which would be mean, standard deviations, or correlations in this unusual model. Given quite small differences in those inputs it produces extreme portfolio weights.

Rob Carver: 00:12:02 The second thing is that it assumes that those inputs are known with certainty, and as we've already discussed, there's very little in finance that we know with certainty. And trying to forecast these things into the future is even harder. There's been a sort of cottage industry that's grown up around making these optimization techniques better, with respect to those two problems. But that comes at a cost, and the cost it comes at is a real loss of intuition and transparency to how the optimization is actually producing its results.

Rob Carver: 00:12:40 One thing I noticed when I was working in the industry was that we typically use one of these techniques to actually work out what our weights should have been in the past, when we're doing a back-test. Which is a sensible thing to do, because a back-test ideally should sort of put you back into history and make sure you only make decisions based on the information you had at the time. The only way of doing that is to use a method that can be automated, like these complex optimization techniques can be.

Rob Carver: 00:13:20 But then when it came to actually choosing the portfolio weights we actually wanted to trade with, we typically use what are called, what you can rather grandly call heuristic methods. But actually they're sort of more like rules of thumb. For example, if you had 10 assets and they're all pretty similar, then it makes complete sense just to split your portfolio 10 ways. If one of those assets was a little bit different, a bit more diversifying, or perhaps you had good reasons for thinking it could have a better performance, then you'd give that a relatively higher allocation, compared to the rest of those assets.

Rob Carver: 00:14:02 Handcrafting was me trying to sort of formalize that process, and write down what you should actually do. The advantage of using this kind of approach is that you could put together a way to optimize portfolios that you can sort of do, I like to say

jokingly, on the back of an envelope. But, in practice you'd probably need a spreadsheet, unless you've just got a very small number of assets.

Rob Carver: 00:14:29 But the other advantage of that is that, because it is using heuristics that humans find intuitively acceptable, you can kind of explain to people how it's working, and explain the components. Well, the reason it's giving a higher weight here is because this asset has better than expected Sharpe ratio, or has lower volatility, or is more diversifying. And you can do that in a very transparent way.

Rob Carver: 00:15:00 You kind of reclaim the transparency and intuitiveness that a good portfolio optimization process should have. Then you can create the implementation in such a way that it allows for the uncertainty and the estimates, and it also will not create portfolios that are extreme, because it's always being pulled back to the idea that unless you have a good reason to do otherwise, you should just allocate your portfolio weights equally.

Adam Butler: 00:15:35 Right, and I remember in your most recent book you spent quite a lot of time on describing a process for handcrafting a more strategic portfolio for a long-term policy asset allocation. It seemed like you applied this sequentially more granular, hierarchical approach. Can you dig a little deeper into how that works?

Rob Carver: 00:16:08 Yeah, one of the ways to make portfolio optimization more tractable for human beings is to do it in stages. If you were to do a strategic asset allocation, you'd want to do that in a top-down fashion. The first thing you should do is say, "Well, I'm going to allocate between asset classes." Or, perhaps even between like a soup of asset classes. You might want to, for example, lump bonds and other products that are ... interest rates, related together and lump equities with like volatility products, like VIX ETFs for example.

Rob Carver: 00:16:50 But you start at the asset class level. That gives you a relatively small number of things you've got to be worrying about. You might have two, or three, or maybe four elements in your portfolio you're trying to find the correct weights for. Then once you've established what the weights should be at that level, then you dive down into each of those asset classes and say, "Well, what should the weights be within the asset class, within that element?"

Rob Carver: 00:17:18 For equities, you might say, "Well, the next stage is I'm going to decide what my weight should be to, say, developed and emerging markets." Then you go down another level, so within developed markets you then get into regional allocations. This continues down to sectors, and then eventually you'll be at the point of individual shares, where obviously that's kind of the atom in portfolio construction. You can't go any more granular than that.

Rob Carver: 00:17:47 This also tends to produce more robust results, so some of the more complicated portfolio optimization techniques that I've talked about. There is one called Hierarchical Risk Parity, or HRP that breaks the portfolio down in exactly the same

way, and then does the optimization of the individual elements, of which at that point there are only a small number. We know kind of theoretically that that tends to produce better results.

Rob Carver: 00:18:23 You also get similar results if, for example, you're trying to combine forecasts together to predict something. If you do that hierarchically you get better results, you get better combined forecasts. It's generally the case that putting things in hierarchies tends to produce better results. But also it means that the problem becomes more tractable for a human being, either because they're doing the optimization themselves, which perhaps you wouldn't normally do, but alternatively the computer's done the optimization, you want to understand how it's done it. If you just got a look at, at most, a handful of assets then it's much easier to understand that.

Adam Butler: 00:19:05 Yeah, that makes sense. One of the ... well, there's a few fundamental assumptions that you make in order to construct portfolios like this without forecasting returns. What goes into that?

Rob Carver: 00:19:27 Let's just be clear, generally speaking I believe that it's pretty easy to forecast standard deviations, which means, generally speaking, you should be using standard deviations as an input into your portfolio optimization. I also believe that it's pretty easy to forecast correlations, although not quite as easy as for standard deviations. That means, again, you should be bringing those into your optimization, although slightly pulling again towards equal weights, because correlations aren't known with such a degree of certainty.

Rob Carver: 00:20:04 The thing I kind of have a problem with, is the idea that you can confidently forecast Sharpe ratios, so risk adjusted returns. To be pedantic, there's a difference between trying to forecast a return and trying to forecast a Sharpe ratio. Generally speaking, if you think you can forecast risk pretty well, and you then say, "Well, I can't forecast Sharpe ratios, I'm going assume that everything has the same Sharpe ratio," then implicitly you are going to produce a series of return forecasts. But that's just going to be your risk, which you think you can estimate pretty well, multiplied by some assumption of the constant Sharpe ratio.

Rob Carver: 00:20:44 The reason I'm generally uncomfortable with the idea of predicting Sharpe ratios, certainly as a starting point, is that it's really hard to do. If you look at the concept of uncertainty measured as, in statistical terms, as the uncertainty of in-estimates, which in layman's terms is basically saying, "How accurately did we know the estimate of this thing looking at past data?" So, not even trying to do a forecast.

Rob Carver: 00:21:20 Looking at the amount of error you get when trying to estimate standard deviations is pretty small. Correlations, it's a bit bigger. But with Sharpe ratios it's really large. For example, if you were to have, say, 20 years of data, and a pretty decent Sharpe ratio of, say, .5, you still wouldn't have enough evidence to say

statistically that that was a positive Sharpe ratio. It's the most difficult thing to forecast.

Rob Carver: 00:21:57 It also has the largest effect on the weight of your portfolio. Putting in a Sharpe ratio estimate that's just slightly different in our naïve optimization, will be the thing that sends your portfolio to extreme weights and allocating everything to the asset that looks slightly better. The consequences of trying to forecast Sharpe ratios with a normal optimization could be pretty bad, I would say, and lead to pretty poor outcomes in terms of the portfolios that you get.

Rob Carver: 00:22:33 Now, that's not to say that you shouldn't try and forecast Sharpe ratios. Anyone who does anything other than buy and hold investing, they're probably going to be doing some kind of Sharpe ratio forecast, right? As someone who has a lot of trend following in their system, I'm saying that I believe that Sharpe ratios can be forecast, conditional on whether an asset has moved up or down in price recently.

Rob Carver: 00:23:00 Now, the question then becomes what you actually do with that information. I've got a kind of classic CTA futures strategy, and in that strategy the right thing to do is to basically just use the trend as it is, the trend signal, and translate that directly into a position. But in the part of my portfolio where I'm doing the strategic top-down asset allocation, and then saying, 'Well, I'm going to make tactical bets on top of that,' based on, say, momentum. In that case, the portfolio is only tilted a little bit one way or the other, depending on what the momentum signal is saying. That's a reflection of the fact that, knowing or having a prediction for the Sharpe ratio changes the expected distribution of returns.

Rob Carver: 00:24:00 But it doesn't give you enough information with little uncertainty that you can say, "Yep, you know what? The last 12 months equities have been up, therefore I'm going to put 100% of my portfolio in equities." The signal just is not strong enough to do that.

Adam Butler: 00:24:17 Yeah, in my experience I wonder what your experience has been, but when speaking with investors who are just beginning this journey, many start with investigating strategies that are all in or all out, so completely binary. Stocks are above the 200-day moving average, so I'm in, and stocks are .001% below the 200-day moving average at the end of the month, so I'm going to be all out.

Adam Butler: 00:24:48 Do you observe this, that a lot of people sort of start there with this 'you're 100% certain in one condition, 100% certain of the opposite effect in the other condition,' or ... how did your thinking on that evolve through time?

Rob Carver: 00:25:05 Yeah, I think it's just natural, because human beings aren't very good at seeing shades of gray. We've only evolved really over the last five, 10, 15,000 years. Before that we were in an environment where you couldn't sort of go, 'Oh, look there appears to be a dangerous animal running at me across the plain. I'm just

going to do a Bayesian optimization to work out the probability that that is in fact going to eat me.' You just have to make an immediate decision to fight or flee.

Rob Carver: 00:25:35 Our brains aren't great coping with uncertainty with shades of gray. Now, I think that there's a difference between a strategy which you use as part of your portfolio. You might, as an investor, say, "Well, I'm going to have a core portfolio with a long only asset allocation to the usual asset classes. But then I'm going to put maybe 5 or 10 percent of my portfolio into a trend following fund. Now, in that trend following fund, they may well be shifting entirely from 100% long S&P, to 100% short S&P, depending on the trend following signal, but that's only a small part of your portfolio. It's acting as a diversifier.

Rob Carver: 00:26:19 Now that's going to be, in terms of outcomes and market risk, that may well be functionally equivalent to a portfolio where you were doing the market timing yourself, but you're just tilting. Instead of having a 5% allocation to the on/off world, you say, 'Instead I'm going to allow my own weights to move, within a 5 or a 10 percent band.' Maybe you'll go from 60:40 equities/bonds to 70:30, or possibly down to 50:50.

Rob Carver: 00:26:54 I think investors need to understand the difference in context between those two places. Just because you read in the newspaper that some high flying global macro hedge fund trader has gone massively short the S&P index, doesn't mean that you should do the same, because the mandate they have is to achieve absolute return, and therefore they need to do that.

Rob Carver: 00:27:23 If they turn around to their investors and say, 'Well, we're deeply bearing and pessimistic about stocks this year, so we're going to reduce our allocation from 60% to 50%.' They'd be like, 'Well, guys, we're not paying you for that. We're paying you fees to do the absolute return stuff, we've already got our long-only beta portfolio over here that we're paying much lower fees on.' It's not your job to replicate part of that.

Adam Butler: 00:27:50 Right, and of course investors usually don't know what else is being held by some of these big hedge funds. There may be news that a large hedge fund has taken a large short position in ES futures, but maybe they've also got a large short … position on somewhere, or some other offsetting position that they're using as an arbitrage.

Adam Butler: 00:28:14 There's lots of moving parts, and I don't think the human brain is well-equipped to think about things at a portfolio level. We sort of, as you say, instinctively think about things as being on or off in isolation, but the portfolio provides much broader context that it's very difficult to visualize or to intuit, which makes this very challenging.

Adam Butler: 00:28:41 Just in terms of just digging a little bit deeper into handcrafting, I got the sense from your handcrafting series that, to a reasonable extent, it relies on very long-

term estimates for volatilities and correlations, and some sort of broad assumption that they're relatively stable through time. It reminded me a little of that cartoon of the economist who put his head in the oven and his feet in the freezer, and saying, "Oh, in the middle or on average I'm just right."

Adam Butler: 00:29:16 Do you make adjustments? How frequently? We know that volatility and correlations can be highly unstable. Even over long periods I'm just thinking back to the 1970s where stocks and bonds both had a very high positive correlation on the order of .5 or .6, contrast that with the recent decade where they've had a reasonably negative correlation. I mean that can have pretty profound impacts on optimal portfolio construction.

Adam Butler: 00:29:43 How do you think about that, these estimates changing through time? And how should investors manage that?

Rob Carver: 00:29:50 Again, it depends a lot on the context. Let's start with assuming that we're looking at individual instruments, so we're allocating to stocks or maybe to ETFs, but not to trading strategies. In that context, pretty much the best estimate of what volatility's going to be in the future is something around a 30-day estimate of volatility. That's what risk metrics use. I use an exponentially weighted moving average … with the same half-life, roughly.

Rob Carver: 00:30:20 Anything up to about six months back works pretty well. Beyond that, you're right, it absolutely becomes more unstable. In the case of some asset classes like equities, obviously you could be talking about perhaps a four or even a fivefold range in terms of [volatility estimates, which is going to make quite a big difference. For volativity I'd use quite a short lookback. Correlations, I'd probably use a longer lookback.

Rob Carver: 00:30:54 I find that the sweet spot is somewhere between maybe six months or a year perhaps, looking backwards. Less than that, and there's just too much noise, although the handcrafting method is designed to be robust to correlations moving around a bit. Obviously, you can get sort of brief shocks that temporarily make markets move together when they don't normally, or vice versa. And generally speaking, you want to avoid those, because trying to trade around them will give you a lot of extra trading costs with no additional benefit.

Rob Carver: 00:31:39 Now, for Sharpe ratios, I guess it depends on what you're using to forecast them. The sweet spot for trend following is a lookback of somewhere between three months and 12 months, depending on the asset class. That's how fast your Sharpe ratio estimates will be changing. Over a year in a steep sell-off you might expect your equity allocation to move from kind of one end of your limit to the other, quite plausibly.

Rob Carver: 00:32:17 Now, the picture changes if we're allocating amongst trading strategies rather than instruments. The first thing is that the volatility generally is something that you

could ignore, because at least with futures, any product where leverage is, to a first or an approximation, freely available. You can scale the individual elements of your portfolio. They've all got the same expected volatility, so we could ignore that.

Rob Carver: 00:32:49 The correlation between trading strategies is both lower and more stable than it is for individual instruments, generally speaking. I'm actually happy to use a … lookback. I use about five years, actually, looking at trading strategies. The Sharpe ratio ... so this is now getting a bit meta. Meta, meta, meta, again. We're trying to predict the performance of something that in turn is trying to predict the performance of something else, but anyway. I just generally don't do that. I just pretty much assume that my Sharpe ratios for trading strategies, in expectation, are all pretty much the same, because, and this is something that Cliff Asness writes about quite a lot. In general, factor timing is difficult.

Rob Carver: 00:33:41 In other words, dynamically trying to allocate between different sources of return like value or momentum, is a difficult thing to do. It's really, really hard to do. So, well, actually I'm not going to bother trying to do that, which means essentially that I assume that all my trading strategies have the same expected Sharpe ratio.

Adam Butler: 00:34:06 It sounds like there are, depending on the strategy that you're running, obviously there are different ... you'd have different guidance, in terms of how quickly you would want to rebalance the frequency at which you ... and lookbacks that you use to estimate the different parameters that you're using to form the portfolio.

Adam Butler: 00:34:32 Do you advocate for a more dynamic, robust, automated optimization approach for your more sort of faster-moving strategies? Or, do you still try to apply a more manual, kind of handcrafting, approach to that?

Rob Carver: 00:34:51 Do you mean in terms of allocating across sort of fast-moving strategies?

Adam Butler: 00:34:56 Well, yeah just your trend following strategies, your carry strategies, that sort of thing, where everything is automated, you're able to observe volatility and correlations in real time, and apply robust optimization methods to those.

Rob Carver: 00:35:12 Yeah, in the trading strategy world volatility scaling, as I would call it, the world's constantly adjusting position to achieve a given target risk, where that risk, in turn, is coming from a risk allocation, a risk budget, and a concept of how strong the forecast or the signal is. That's happening continuously. I pretty much ... you could say I trade on a daily basis with my future systems. I'm effectively updating my forecast of volatility every day, and if necessary rescaling my position to adjust for that, again, on a daily basis if necessary.

Rob Carver: 00:36:02 The actual risk budget then that is given to each individual trading strategy, where a trading strategy could be something like a moving average on a particular instrument like the S&P 500 future, that is a much more slower-moving thing.

Because, what evidence would you have for updating those numbers? It would be because either the correlation of the individual trading strategies has changed, or because the Sharpe ratio, you're conditional expectations of Sharpe ratio, have changed. I've already said that I don't believe that the last ... the second of those, the Sharpe ratio, is something that's worth considering.

Rob Carver: 00:36:47 Then we have to say, ' Well, given that I've got a five-year lookback on these strategy correlations, actually, on a month-to-month or even a year-to-year basis, those risk weights probably aren't going to be moving very much.' In practice I've actually kept the risk weights, so my own system, constant for the last five years. Although when I was working in an institutional environment, those risk weights would change generally when one or two things happened. Either they would change annually, or much more likely, in between those cycles there would be some change to the system, so some innovation, like a new kind of signal, or a new market being traded.

Rob Carver: 00:37:31 My research cycle is much, much slower than that. I'm actually keeping those weights as good as fixed, effectively.

Adam Butler: 00:37:41 Sorry Rob, I've got a technician at my door who was supposed to be here two hours from now. Can you give me just like 30 seconds to let him in?

Rob Carver: 00:37:49 Yep, that's fine.

Adam Butler: 00:37:49 Thanks.

Adam Butler: 00:38:55 Hey there.

Rob Carver: 00:38:56 Yep.

Adam Butler: 00:38:57 You go home or where you could hope to get some quiet, and then, interruptions everywhere.

Adam Butler: 00:39:06 Got it, so we were talking about the fact that you keep your allocations fairly constant to the actual trading strategies, because the correlations don't change very much, and you're managing the volatility of each of the strategies directly. If you don't have any cause to believe that there's been a major shift in expected Sharpe ratios, then there's no need to make any meaningful adjustments to the allocations. Did I summarize that reasonably?

Rob Carver: 00:39:43 That's correct, yeah.

Private Investment Thoughts

Adam Butler: 00:39:45 Perfect. Just shifting gears a little bit, and thinking about the typical kind of institutional, or family office portfolio and how that's evolved over the last 10 or

15 years, I'm noting a large proportion of assets flowing towards private instruments, private equity, infrastructure, that sort of thing. Are you observing a similar flow? And what are your thoughts on private investments?

Rob Carver: 00:40:19 Well, I'm probably not in the best position to judge that, because I've only kind of got anecdotal evidence now, not working in industry, that I kind of gather when I do what I tell my wife, ' intelligence gathering ' trips to the nearest local pub. But, it's interesting because actually this is like history repeating itself, because if you go back sort of 15 years, you will probably remember there was a big fuss about the fact that the big US university endowments, your Harvard's and Yale's, were making a big deal about going into alternative assets and in particular hedge funds, but also private assets.

Rob Carver: 00:41:04 Now, as an economist, you'd expect that. The return from private assets would be higher than the return from public assets, because private assets have obvious disadvantages. They're harder to trade; they're less liquid; there's a lot more information asymmetry, the kind of legal and regulatory overhead in terms of thinking that you're getting a true and fair financial statement, is not as good for private companies.

Rob Carver: 00:41:40 Although, obviously we could both quote plenty of examples of public companies which you know went spectacularly bust with accounting fraud happening somewhere. We have recent examples in this country actually. I guess as a quant, I'm uncomfortable with private assets, because the data's not there. It's much harder to get a time series of prices for a private asset. And even if I can get some kind of time series, there will be a lot of uncertainty about each of those data points.

Rob Carver: 00:42:21 It might be that, yeah sure, on these given days, some VC valuation was put in that theoretically gave the company this valuation per share. But what's happening in between those times, and whether that really was a true market price, you can't say. So I'm naturally biased against private assets from that perspective.

Rob Carver: 00:42:50 Diversification is a thing you should have in any portfolio. Probably at least 40% of kinds of investible assets, say in the UK and the US, are private. If you don't allocate to private assets, then you're kind of missing a big chunk of what's out there, you're missing a big chunk of diversification. Then secondly, from a theoretical perspective, investing in private assets should give you a higher return, because it's going to give you potential problems. You should be rewarded for those problems.

Rob Carver: 00:43:29 I don't see any reason why someone shouldn't put some of their portfolio into private assets. I guess the issue then becomes, well, can you get adequate diversification? As a private individual, let's suppose I wanted to put, I don't know, $100,000 into private assets? Well, perhaps I can put $50,000 into two companies, because a smaller ticket size in that, is just not going to be economical in terms of

the legal fees I'm going to have to pay, in terms of due diligence I'm going to have to do. Whereas $100,000 in public companies, well I could just buy an ETF, I could buy MSCI Global ETF, and effectively have exposure to, I don't know, seven and a half thousand companies worldwide.

Rob Carver: 00:44:21 What's the minimum number of companies for an adequately diversified portfolio of private companies? Is it 10? Is it 50? Is it 100? I don't know, but it's probably going to be at least 50 or 100, in which case you're talking about pretty large family offices. They're going to be at the point where they can actually say, "Well, yeah, we're going to put 10% of our portfolio into private companies, and then we're going to split that amongst 100. And we're going to have to recruit like 25 people ..."

Adam Butler: 00:44:58 Right.

Rob Carver: 00:44:59 "... to oversee that part of our portfolio," whereas it could be that a large chunk of that portfolio is in long-only beta, which requires like one guy to check it every month, and maybe 10% is in hedge funds, which requires maybe five people to do due diligence for.

Rob Carver: 00:45:22 I think the standard thing in finance and economics is there is no free lunch. Sure, private assets probably give you a higher return, but at some cost.

Adam Butler: 00:45:32 Yeah, I mean the irony is that I struggle to find any literature that validates the view that they actually produce a higher return. Certainly a higher return than you could get by simply levering a mid-cap equity portfolio, and at least you'd have visibility, transparency, and liquidity. But after you factor in search costs, financing costs, deal costs…

Rob Carver: 00:45:59 Yeah, last year there was some research done by somebody looking at private equity funds. I can't remember who did it, but they're basically all the ... the private equity funds always put up these great charts showing how well they're investing their clients' money. But it turns out that actually those are based on some really fundamentally flawed assumptions, and that actually the average investor in a private equity fund has lost money.

Rob Carver: 00:46:34 There was obviously going to be a problem getting this research done, because the data is going to be very hard to find.

Adam Butler: 00:46:40 Yep.

Rob Carver: 00:46:43 It would be an extremely costly and expensive exercise to do. We are kind of stuck with this theoretical and intuitive feeling that yeah, we really ought to be paid more to invest in private assets. And quite a lot of anecdotal evidence of people who have said, ' Yeah, I invested in,' I don't know ' a chain of laundromats, and

now I'm a millionaire.' But, I know a few counter examples in the other direction, but-

Adam Butler: 00:47:09 Yeah…

Rob Carver: 00:47:10 There's a quant among … for a lack of hard evidence, definitely.

Adam Butler: 00:47:14 Yeah, I personally think that the primary attraction is the autocorrelation of returns, so you just don't see most of the volatility. But, of course, unfortunately that means that the volatility is being mis-measured. You're loading up on assets that are fundamentally designed to have strong procyclical risk characteristics, so are not at all the portfolio diversifier that you think that they are, especially when you need them to be a diversifier.

Adam Butler: 00:47:47 I think it's a very misunderstood asset class, and very misunderstood investments, and I struggle with it.

Rob Carver: 00:47:55 Yeah, I'll just make one more quick point, which is that you can think of investing in a, I don't know, let's pick a … fashion example. Investing in General Electric is essentially the same as investing in about 5,000 small private firms across a diversified group of businesses.

Adam Butler: 00:48:13 Right.

Rob Carver: 00:48:17 If we then say, ' Well, actually what we can do, you can kind of back out the expected volatility of an S&P 500 stock by looking at the volatility, the index, and the average correlation of those stocks.' If you then do a similar exercise with those, it will tell you that the expected volatility of those individual private firms is going to be two or three times the volatility of General Electric.

Rob Carver: 00:48:44 It's going to be a pretty high number, it's going to be quite a high volatility.

Adam Butler: 00:48:49 Mm-hmm (affirmative)

Rob Carver: 00:48:51 As you say, it's hard to measure in practice. The way that it's measured dampens down the volatility. You get a similar effect if you look at house prices, and this is versus the price of an individual house. But I think that the risk of these investments is very much underestimated by a lot of people, definitely.

Adam Butler: 00:49:13 … two quant's views on private investments. Just broadening this subject a little bit, what do you see as some of the major missed opportunities for many institutional investors and large family offices, in terms of strategy allocations, or just general investment theses?

Missed Opportunities

Rob Carver: 00:49:35 Most institutions tend to be a little bit kind of followers of fashion. At any given time, they are either under or over allocated to a particular theme or a particular style. In some cases they've got no allocation at all. Based on our conversation just now, it strikes me that the average family office has probably got too much allocation to private assets.

Rob Carver: 00:50:15 The question then becomes, what is it in their portfolio that they've got less of an allocation to? The way that these guys tend to invest is countercyclically, so they will get out of things that have done badly in the last year. Now, I'm a big fan of trend following, but trend following tends not to work very well when you're talking about allocating to investment styles.

Rob Carver: 00:50:40 I said earlier that I believe that trying to time factor returns is very difficult. But, trying to time factor returns by getting out of the factor or the theme or the style that did really badly last year, is even worse than-

Adam Butler: 00:50:56 Especially toxic, yeah.

Rob Carver: 00:50:58 Yeah, so most likely, investors will not be in trend following. Most likely they will be well out of emerging market bonds and emerging market equities. Your typical large family office or institutional portfolio probably isn't missing a trick in terms of the fact that there are things out there that they've got no allocation to at all. But they are missing a trick in that their allocation policy is too time varying and inconsistent, and has no systematic underpinning, which means that they will be underweight a bunch of things, which they probably ought to have a higher weight to … because it will be giving them a more diversified source of return.

Rob Carver: 00:51:49 Theoretically, we should try and diversify our portfolio amongst different asset classes, and then stage 2 is to say, ‘ Well, actually where does return come from? It comes from risk factors or return factors. So, therefore we should try and diversify ourselves across different return or risk factors.’ Which means, for example, that if you're already heavily invested in equities, then dumping a pile of money into a short volatility ETF is not going to give you an extra diversification, because it's the same source of return. It's the equity premium except, because it's an inverse volatility ETF, it's the equity premium squared or cubed.

Rob Carver: 00:52:35 I think ... part of the problem is that a lot of the literature in this area is quite old now, and it's concerned with things like inflation.

Adam Butler: 00:52:49 Right.

Rob Carver: 00:52:49 And how the idea that equity and bonds are both driven by a few different factors, including interest rate risk, including inflation, including earnings growth. Inflation is this sort of thing that's not around at the moment. I think the sort of intuition about the idea that these kind of latent ... what I call, a very fancy way of putting it, but these ... they are latent factors, they are things you cannot see or observe.

Rob Carver: 00:53:27 You can try and come up with proxies for them, but the idea that there should be more than one source of return out there I think has been lost a little in this world of inflation, because that was quite an intuitive way of explaining it to people, to say, ‘ Well, yeah actually, the kind of main risks to the economy, how interest rates move, inflation moves, economic growth.’ You want to kind of be hedged against states of the world when inflation's rising or interest rates are falling, and so on and so forth.

Rob Carver: 00:54:00 With inflation going away, I think people will become much more focused around the idea of there being this one big risk factor out there. What that risk factor is, is the flavor of the month, it changes every week. It could be Trump, it could be China, it could be Brexit. They're very concerned about trying to get that portfolio right, for that kind of one axis. When ... it's actually the world is more complicated than certain than that. There are a lot of axes, a lot of different sources of risk. You've got to try and do the best job you can of diversifying your portfolio across all of those.

Adam Butler: 00:54:39 It really is astonishing to me. I have conversations with really bright, typically more sort of systematically minded investors, just because of the way that we approach the problem. There seems to be this fairly pervasive view, that a maximally diversified, some kind of risk parity, or global risk parity type strategy, however you want to construct it, the handcrafted strategic portfolio that you describe in your portfolio, I would characterize that as a sort of heuristic, global risk parity strategy.

Adam Butler: 00:55:21 We sit down with guys from AQR, quants that don't run risk parity strategies. My observation is that everybody seems to feel like this is probably the most legitimate, least biased place to start for portfolio construction before you begin to introduce active bets. And yet, when we go out and talk to institutions, or try to put this in practice, there's just almost no uptake. Where's the gap there, do you think?

Rob Carver: 00:55:55 Well, I mean I could ask you and say what objections do they raise to it?

Adam Butler: 00:56:00 Well, I mean some of it's policy portfolio. But to my mind that just speaks to the fact that the policy portfolio has been mis -specified.

Rob Carver: 00:56:11 Yeah.

Adam Butler: 00:56:12 Right? If your answer to the question ‘ Why shouldn't I do something?’ is ‘ Well, because we haven't done it before,’ then probably you're missing some pretty substantial opportunity.

Rob Carver: 00:56:25 I do remember that, this was a few years ago, there was a lot of criticism about risk parity, and I guess in particular leveraged risk parity. The idea that you could say, ‘ Well, yeah the optimal portfolio, so let's just say bonds and equities, inverse

volatility to assets,' inverse volatility is the same as risk parity, ' you're probably going to want something like 70:30 bonds/equities, roughly speaking in terms of cash flows.'

Adam Butler: 00:56:55 Yeah.

Rob Carver: 00:56:56 And that's going to give you a really low expected standard deviation of perhaps 3, or 4, or maybe 5 percent a year, tops. Therefore, assuming a Sharpe of say, a half, you're going to get 2 and a half percent a year, plus risk-free another couple of percent. For most people that's obviously way too low. So, then you add leverage. You say, ' Well, I'm going to leverage this thing two or three times.'

Rob Carver: 00:57:23 I think when people look at that kind of portfolio they raise a number of objections, and the problem is they're just pointing at aspects of it and saying, ' Well, I don't like this, I don't like this, I don't like that, I don't like that.' They're kind of confounding all of their objections together. Some people are naturally uncomfortable with the idea of leverage, which is not unreasonable.

Adam Butler: 00:57:46 Except, I'm sure they borrowed to buy their home-

Rob Carver: 00:57:49 Well, exactly, and-

Adam Butler: 00:57:49 If they own any S&P 500 they've got a 2:1 debt to equity ratio, on average.

Rob Carver: 00:57:57 Exactly. But there's a big difference between, actually the home buying thing is exactly the same thing, but there's a difference between latent or implicit leverage, it's actually you that could be getting a call from your broker saying, ' I need more margin.' So that's fair enough, to a degree.

Rob Carver: 00:58:19 But then they also say things like, ' Well, you know, you're allocating 70% of your portfolio to bonds, and we have just had a 30-year secular bull run in bonds, the Fed's going to start pulling up rates, you know you're basically putting nearly all your eggs in a basket of assets which I know for sure is going to go down in price.'

Rob Carver: 00:58:41 The problem is the reasons why the ... well then, how do you address that? You can say, ' Well, it's fine to not have 70% in bonds, and to get your edge to return not through leverage but by allocating more to equities. So you'll get a lower Sharpe ratio, but you'll get a higher expected return.' Then you can … to say, ' Well, you shouldn't put all of your money in equities.' One of the reasons for that is because of the difference between arithmetic and geometric returns - you then enter quite a long and complicated discussion into explaining why this thing makes sense.

Rob Carver: 00:59:29 Because just the reasons why these portfolios make sense are essentially because you can predict volatility really well, so the idea of weighting inversely to volatility, a risk parity idea, makes perfect sense. You could predict correlations pretty well,

which means you should have a portfolio that's as diversified as possible. But actually you can't predict returns or Sharpe ratios very well.

Rob Carver: 00:59:58 Ultimately, if you want higher returns you're going to have to pay for it with increased risk, and that risk can either come from leveraging up a 70:30 portfolio, or it can come by moving to an 80:20 portfolio, which will probably have an inferior Sharpe ratio - well it would definitely have an inferior Sharpe ratio. It may even have an inferior geometric return, which means that actually, in 50 years or whenever you cash in your funds, the expected final value of that portfolio will be lower.

Rob Carver: 01:00:34 The concepts that stand behind these portfolios meaning that they make sense, are A, unintuitive, and B, quite complicated to explain, whereas the attacks that you can make on them, the criticisms you can make on them, are very intuitive and easy to understand.

Adam Butler: 01:00:49 Yeah, we typically start from the place where you started earlier, which is this idea that you're rewarded for accepting risk on growth, or on inflation, or rates. You want to have markets in your portfolio that are fundamentally designed to hedge against each of those major states, right? So, worse than expected growth but higher than expected inflation, or better than expected growth but higher than expected inflation.

Adam Butler: 01:01:27 Depending on how many ways you want to slice it, there's between four and eight states of the world, and you want to own assets that are designed to hedge against all of those different states, and if you hold them so that they're in proper balance, then your portfolio should be relatively resilient against whatever the market throws at you over the intermediate term.

Adam Butler: 01:01:50 It's funny, because you get all kinds of heads nodding there, and then you slowly sort of build into what the portfolio looks like, and that's the point where-

Rob Carver: 01:02:00 Yeah.

Adam Butler: 01:02:00 ... the objections that you've just described come into play.

Rob Carver: 01:02:03 Yeah.

Adam Butler: 01:02:04 And it seems like that theoretical basis is completely lost as you sort of look at the portfolio weights.

Rob Carver: 01:02:11 Exactly, if you could explain to people ' look, do you not agree that if an asset is risky you should have less of it in your portfolio?' And yes, they nod. ' Do you not agree that a portfolio should be diversified across different asset classes, and countries, and so on?' And yes, they nod. And then you say, ' Well, actually what you should really have is 70:30 portfolio with 2:1 leverage.' And they freak out.

Three Judases

Adam Butler: 01:02:30 Right. Exactly, it's a very strange phenomenon. We face this all the time, and I'm wondering, obviously you've had similar conversations which is encouraging, but also discouraging. Just, I really wanted to cover some of the topics in your most recent presentation which you've called “Three Judases”, which I thought was just really creative, and I'm going to link to this presentation in the show notes because I think it's absolutely fabulous. I want to make sure we do carve off some time to cover it.

Adam Butler: 01:03:02 To start off, I thought you had a really clever way of attacking this idea of the bias variance trade-off, and you sort of reframed it as a trade-off between model complexity and prediction error. Can you just talk for a couple minutes about what that is?

Rob Carver: 01:03:21 Yes, the idea basically is that if you were trying to create a trading strategy or something else that's trying to predict the future, the first thing that you do in a quantitative world is you get some data and you say, ‘ Well, how well can I ... you know, what kind of strategies should I produce? With what parameters that will do a good job of predicting what happened in this particular data set?’

Rob Carver: 01:03:51 When you do that, what you will find is that the more complicated you make the model, the more parameters it has, the more moving parts, the better a job you will do. Actually, there is a point at which you will be able to predict the past perfectly. To take a trivial example, if you took a year's worth of returns, you could create a model that said on January the 2nd the market will go up, on January the 3rd the market will go down, and so on and so forth.

Rob Carver: 01:04:21 You have one of those for every business day of the year, and you'd be able to predict exactly what had happened in the past, because the complexity of the model was exactly enough to do that.

Rob Carver: 01:04:34 Now the issue, of course, comes due to the fact that we do not get paid for predicting the past, not directly at least. Maybe I'll move on to that in a moment, but we only get paid for predicting the future. The way we try and simulate this when we're doing this type of work is by taking another set of data that the algorithm fitting process hasn't seen, which is called the test set, or the out-of-sample is also a common phrase, and then we see how well the algorithmal strategy that we've trained on the original data does in this out-of-sample sample data.

Rob Carver: 01:05:14 This effectively is simulating what happens if you build a trading strategy based on historic data, and then go out and trade it with real money, and obviously when you're trading with real money, that's by definition a data set that the algorithm hadn't seen, because it didn't exist, it's only in the future. Now, what happens here is that, as you add complexity to the model, the out-of-sample performance, to

begin with gets better. So a very simple model which just said, ‘ Go long all the time,’ is going to do okay.

Rob Carver: 01:05:52 In asset classes that obviously have gone up, if you make it a little more interesting and maybe throw in a trend following component, or carry or value or whatever, then it will do a bit better. But at some point the complexity starts to get in the way and the performance of the model, actually on an out-of-sample, degrades. We take the simple example of a strategy that had a different prediction for every day of the year, that probably worked really well in 2018. It's unlikely it will work as well in 2019 - extremely unlikely, in fact.

Rob Carver: 01:06:27 The trick is to try and create a model which has an optimal level of complexity. Not so little complexity that it doesn't actually do anything useful, but not so much complexity that it's no longer robust, and it's too overfitted or curve fitted to the historic data. This is a really hard thing to do, because although I said you don't get paid directly for predicting the past, in fact you kind of do. If you're an academic researcher and you write a paper that finds some interesting market anomaly, and you get it published, you'll get tenure and that will do well from there.

Rob Carver: 01:07:09 If you're working in a quantitative investment firm, and you produce a back-test showing that this strategy has traded successfully, you'll get a nice bonus. And even me just sitting here, if I was to create a back-test with unrealistically good results, I would probably feel good about myself, so there would still be some intangible benefit from doing that.

Rob Carver: 01:07:34 But when it actually comes to ... so there are all of these biases pushing us towards making our models more complex so that we can improve past performance, because that's more tangible and can actually lead to tangible rewards, but more complex models will then do badly out-of-sample when trading with real money, which of course is what people actually care about.

Adam Butler: 01:07:55 Right. How does this fit into your framework of explicit, implicit, and tacit overfitting?

Rob Carver: 01:08:03 Explicit fitting is kind of what everyone mostly does. Generally speaking, when you're fitting some kind of trading model, you will have an automated fitting algorithm which could be something very simple like a grid search, could be a regression, or maybe you're fitting is a portfolio allocation, so you're trying to decide what weights to put on different predictors, or different factors, or different level trading strategies, in which case maybe you do a market fit optimization, or something else.

Rob Carver: 01:08:38 The key point here is that the fitting is happening automatically, that means that you can run it on a pure out-of-sample basis. You can run the fitting automatically, iterating it forward, ensuring that the fitting the algorithm can only see is data that was historic to before that, so you could do true out-of-sample tests. You can also

control the degree of parameterization, how complicated did you let the model become? There are ways of doing that sort of automatically, where the algorithm can just say, ‘ Well, how complex will the model be, before it gets too complex?’ It's a controlled process.

Rob Carver: 01:09:23 Fitting gets really bad press, and overfitting gets really bad press, but this kind of explicit fitting where it's all being done automatically on a pure out-of-sample basis, that's the best kind of fitting to do.

Adam Butler: 01:09:35 This is your typical sort of walk-forward testing?

Rob Carver: 01:09:39 Yeah, exactly. Yeah.

Adam Butler: 01:09:42 Perfect. And then implicit fitting?

Rob Carver: 01:09:46 Implicit fitting is something that I think intuitively most people in this business have got, everyone's done it. I think everyone knows that they're doing it and they know that they're doing something wrong, but it's really hard to stop yourself. An example of implicit fitting would be this: you create a strategy, and you run a proper walk-forward back-test optimization on it, and you get some results. You look at those results, and you think, 'I think I can improve on this and make it a bit better. I'm going to just change ...'

Rob Carver: 01:10:24 There's kind of different degrees of this. On one extreme would be cheating. Cheating would be to say, actually, the best parameter value over the whole sample set is X, therefore I'm just going to use X for the whole sample set. So that's kind of outright cheating. But you could also, for example, say, "Well, actually I've used a five-year lookback here from … correlation estimates. I've used a five-year lookback to make my correlation. I wonder what happens if I use a one-year correlation lookback." You rerun the thing, it looks better. "Okay, I'll use one-year now." You're still cheating, because you still used the whole data set to make a decision about what the parameters should be. You've still done something that someone who was doing a truly out-of-sample test couldn't possibly do, because they did not have access to all of the data to begin with.

Rob Carver: 01:11:09 That's what I would call kind of meta-cheating. Meta's a prefix I'm using a lot today. My example of meta-cheating might be that you actually changed the way that the fitting algorithm works. For example, you might be using a regression technique, which penalizes large parameter values, where you could change the degree without penalization to improve your results. So, these things are all cheating.

Rob Carver: 01:11:42 Now, there are ways of kind of avoiding this and getting around it. There are corrections you can do, to say, well, on tests you can do to see it's likely that the results you've produced are a bit overfitted through this process. But, ultimately, there is still one thing that pretty much everyone does, which is that once they've tested a trading strategy and fitted it either properly, or maybe a little bit implicit

cheating, they then make a decision whether to actually trade that thing going forward.

Rob Carver: 01:12:19 Most people do not trade bad ideas, they throw them away, and in doing so they have effectively, again, committed this crime of implicit fitting, because the decision to not trade a trading strategy made with a back-test based on the past, that's not something that you could really have done in practice.

Rob Carver: 01:12:41 It might be that you don't actually trade your strategy in light trading, but a bad strategy really ought to be kept into a back-test, because otherwise your back-test is only going to consist of ideas that worked, which means that your back-test results are going to be inflated. That's kind of-

Adam Butler: 01:13:03 This is the whole idea that if you run 1,000 back-tests, that just purely by random chance, 50 of them are going to look pretty good, right?

Rob Carver: 01:13:11 Yes, exactly. Most people are aware that they shouldn't; for example when they sit down at a computer to try a new idea out, that if they do, then do 1,000 iterations to improve that. They know they're doing something wrong. Most people know enough. But even someone who's very careful, and doesn't do that, they'll still reach a point where they decide, 'actually, I'm going to proceed with this idea and trade it with real money or I'm going to drop it.'

Rob Carver: 01:13:38 It might be that your live trading ends up exactly the same, but your stored back-tests will be inflated, which means for example, that you might shoot for a higher leverage target than you should. It might mean that you pay higher costs than you should. But also, if you have investors who are gullible enough to be convinced by back-tested pitch decks, then you'll potentially be fooling investors. And as a result, any investor who's worth their salt will automatically look at a back-tested account curve and discount it massively.

Adam Butler: 01:14:13 Right.

Rob Carver: 01:14:14 … divide the results by two or by four, because they know it's most likely that, even the most honest and good person who's done all the right stuff will still drop a bad idea before they actually let it go into their back-test.

Fund Level Back-test Strategies

Adam Butler: 01:14:29 One of the things you mentioned in your presentation, which caught my eye, one of the solutions to this challenge of implicit fitting, you write ‘ Keep all tested strategies in "a fund level back-test” and use portfolio allocation, apply selection algorithm.’ What does that mean?

Rob Carver: 01:14:52 The idea here is that if I was running a hedge fund and I was worried about this problem, what I would do is every time anyone had any kind of idea, I would

basically keep a copy of that back-test in some central computer system, and then when it came to saying what are the back-test results for our fund? Which is asking the question of ourselves,' What money could we have made in the past, given the access to the historic data that we had at that time?'

Rob Carver: 01:15:22 That might mean, for example, there's a strategy that was doing really well up till about 10 years ago, and you keep that in your fund level back-test. And the fund level back-test would then allocate between your strategies, and it would look at that strategy and say, ' Well, that strategy looks pretty good. We're going to give that a pretty decent allocation.'

Rob Carver: 01:15:40 But then as the performance of the strategy starts to degrade, then your allocation process will automatically reduce the weight given to that strategy, and it might be that in the present day the allocation to that strategy would be so low that you'd actually just exclude it, and this is what I mean by a selection algorithm. You might, let's say you've got, I don't know, 50 different trading strategies. It might be that one of them ends up with a weight of less than, I don't know, .1%, you say, ' Well, it's not worth the time and the effort to run and monitor this thing. We'll go and actually trade.'

Rob Carver: 01:16:17 You're basically simulating the real thing that would have happened, which is that someone had developed a strategy, it would work well for a few years and would have gotten a higher allocation in your live trading portfolio. It would have then done badly, you would have reduced its allocation, and at some point you would have turned it off.

Rob Carver: 01:16:31 The alternative way you'd look at the whole back-test of that strategy and say, ' Well, this thing, you know, actually doesn't end up making any money. We're just going to delete the back-test file, put the documents in the bin, and forget we ever did this.' That means that your historic back-test, what I call your fund level back-test, which doesn't include that trading idea, that strategy, will be inflated. Because it's for the last 10 years, because it won't have the reducing allocation to the thing that was doing badly.

Adam Butler: 01:17:01 That makes sense, and you also mentioned just along the same lines that context matters. For example, high-frequency strategies, you've got more data, and therefore you can make decisions about whether or not an edge may have substantially degraded much more quickly. But by the same token, that means that those edges will probably degrade more quickly, versus a lower frequency strategy like trend following, carry, et cetera.

Adam Butler: 01:17:31 You've got less data in terms of number of trades with which to make a real-time decision about whether or not the edge has gone away, but because of that you can't really ... it's much more difficult for that edge to completely decay away. It's less likely to degrade. Can you say more about that?

Rob Carver: 01:17:55 Yes. The first thing to say is that the higher the performance of, you know, so economically, and if you're making a lot of profit in a particular market it's likely that competitors will move in, unless there are regulatory reasons why they can't, like patterns and so on. You can think of trading strategies in a similar way, so if you find a high-frequency, and it's likely to be a high-frequency trading strategy, but if you find a trading strategy that has a Sharpe ratio of 5, it's unlikely that you'll be able to keep that a secret for very long.

Rob Carver: 01:18:31 Other people will discover it. It may be a short-term effect that's leveraging off a particular thing in the market that's going on. It cannot make sense that someone can find a trading strategy with a Sharpe ratio of 5, and keep it to themselves and keep trading it forever. People-

Adam Butler: 01:18:50 They would end up earning the market very quickly.

Rob Carver: 01:18:51 Well exactly, yeah. Essentially, pure arbitrage opportunities that have a Sharpe ratio of infinity, if you like, are extremely rare and get arbitraged out very, very quickly. Things that are slightly less than pure arbitrage get arbitraged out, … more slowly but still pretty quickly.

Rob Carver: 01:19:12 On the other end you've got strategies like, say, trend following, which have much lower Sharpe's. And one consequence of those lower Sharpe's is that there are long periods of time when they're in draw down, and people like write newspaper articles saying, ' Trend following is dead,' and things like that, or 'Value is dead.' That's another common one.

Rob Carver: 01:19:34 Now, it's less likely that these things will go away. Because, well first of all, the reasons why they're happening are probably down to things like cognitive biases in the human brain, blah, blah, blah. Things that aren't likely to go away anytime soon. But also, because the returns aren't that great, not that many people are going to be interested in buying into these strategies, so it's unlikely that they'll get crowded very quickly and arbitraged out.

Rob Carver: 01:20:07 From a kind of market efficiency point of view, generally speaking there's a continuum between strategies that probably have a higher frequency, probably have quite short holding periods, probably have high Sharpe ratios, but which will be kind of found out quite quickly. You can trade these strategies only for six months before the effect disappears.

Adam Butler: 01:20:33 Or with a small amount of capital, right?

Rob Carver: 01:20:35 And probably with a relatively small amount of capital. And on the other extreme you've got strategies which trade pretty slowly, are based on cognitive biases which probably won't go away anytime soon, you can test them for hundreds of years and the effects seem to be pretty consistent. They're lower Sharpe ratio, they're probably not going to get crowded out.

Adam Butler: 01:20:55 Yeah.

Rob Carver: 01:20:56 Now, when you then sort of say, ‘ Well, okay. How can I make decisions about how to allocate to these strategies, or how can I make decisions about when the performance of these things is degrading?’ If you look at the sort of statistical properties of the estimate of a Sharpe ratio, the higher a Sharpe ratio is, the more certain you can be about it. If your Sharpe ratio is really low, you're … to have enough statistical evidence to say whether Sharpe ratio can be positive.

Rob Carver: 01:21:30 You need a lot of decades worth of data. You may have traded it for decades, but that really means that you can't look at, say, just the last year or three years, or even five years, and say, ‘ You know what, this thing is definitely broken. There's like a 99% chance that the returns over the last five years are negative.’ It's just extremely unlikely that you'll get that kind of statistical confidence, unless the thing has been an absolute dog, and has just really done really badly.

Rob Carver: 01:22:01 But, with high-frequency trading strategies that have much higher Sharpe’s, you've got more trades anyway, so kind of intuitively that means you've got more data points more quickly. So instead of having one trade a day, you're probably going to have thousands or tens of thousands. The Sharpe ratios are higher, so there's much less uncertainty about whether they should be making money or not, and that also means you can spot pretty quickly whether they've degraded in performance. Life cycle of a high-frequency trading strategy might be something like this, and … rises in the market.

Rob Carver: 01:22:34 Within a few weeks or months, you have enough data to confirm that the effect is definitely there, and you don't need that much data because it's such a strong effect. You can trade it for three months, then other people find out about it, the performance starts to degrade, and then you can say, ‘ Actually, looking at the last, even perhaps just the last week of data, it looks like this thing has now gone away, and is no longer there. We can shut that down.’

Rob Carver: 01:23:01 Now if you look at something much slower like trend following, or value, or carry, you've got to expand those time frames massively. It can take decades for it to be clear that an effect exists, and it can take many years or even decades before you have enough statistical evidence to say that it's gone away.

Rob Carver: 01:23:23 I've actually added a recent blog post on this, looking specifically at trend following and basically saying, ‘Actually, you know what? A low trend following hasn't done great in the last few years, since 2014. There's nowhere near enough poor performance, or history of poor performance, to say with any statistical confidence that it's no longer a profitable strategy.’

Rob Carver: 01:23:47 That means, from a value allocation perspective, at best you could tilt your allocation away from something like this very, very, very slightly, like from 50% to 49%, that kind of order of magnitude. But you've got nowhere near enough

evidence to say, ‘ Yeah, we can bring the lever on this thing right down to zero,’ which you could do with a higher frequency, more profitable trading strategy.

Adam Butler: 01:24:14 Yeah, that makes sense. Have you read Andrew Lo's book on the adaptive markets hypothesis?

Rob Carver: 01:24:18 I have, in fact. I'm just seeing if it was on the shelf behind me, so I can prove it, but no, I think it's somewhere on the bookshelf. Yes, I have.

Adam Butler: 01:24:25 Right?

Rob Carver: 01:24:26 Yeah.

Adam Butler: 01:24:27 Because it just rings so consistent with some of the themes that are covered in that book. We've been internally sliding towards the opinion that the secret to sustainable success in markets is finding a very large number of very weak classifiers, or very weak edges, and assembling those thoughtfully. Because the really great edges are going to be arbitraged away very, very quickly. And the only way that you can take advantage of weak edges is if you include many of them in a diversified portfolio of strategies, and it's just not the type of exercise that most investors can successfully complete.

Adam Butler: 01:25:21 I think that's sort of the last bastion of sustainable alpha in markets is systemically finding many, many weak classifiers, or weak features, and then finding smart ways to assemble those in a diversified way. That actually ends up I think being the overall research direction for our team over the next few years, which actually is a good segue, because I was going to ask you, I know you're currently working on a new book. Is that right?

Rob’s New Book

Rob Carver: 01:25:53 Yep, that's correct.

Adam Butler: 01:25:55 What's the new book about? And just along the same lines, any insight on what your research focus is going to be over the next year or so? I would assume it's going to be related to the book, but maybe not.

Rob Carver: 01:26:08 Maybe not, in fact. The book is aimed at more of a retail audience than my first two books. If you think about hard numbers, I think you’ll find in my first book that the minimum example portfolio I showed was something like $100,000 of capital, and in this new book it's $500.

Adam Butler: 01:26:29 Oh, wow.

Rob Carver: 01:26:29 That's the difference. It's focusing specifically on leverage products. That's futures, margin trading stocks, FX, and some things that aren't legal in the US actually, but

that's CFDs and spread bets, but it's basically designed for both the UK and a US audience, so it kind of covers all the bases there.

Rob Carver: 01:26:52 And the reason I chose to focus on the leverage products is that, in terms of doing the most helpfulness to the largest number of people, retail traders generally, where they're getting into trouble, they're getting into trouble with these leverage products.

Rob Carver: 01:27:15 It's kind of not that dissimilar from my first book in terms of its, well here is a systematic way to trade, but there's much more hand holding, but also answering a lot of questions that, creating a lot of fallacies that people have generally, but retail traders particularly. Things like, for example, should you trade a particular market? Are some markets better than others for trading? If you look on the internet, some of the kind of uninformed stuff you get on a lot of trading forums, you'll see things people are saying like, ' Oh, well, you should choose a market that gives you the right vibes,' and all this kind of nonsense.

Rob Carver: 01:28:08 And then I look at the statistical evidence, or show ... explain the statistical evidence, saying, actually going back to the first thing we spoke about, which is uncertainty, if I do a bar chart showing the performance of all the instruments in my own portfolio, it looks like there's a big difference in performance between I think the worst performing market, which I think was the Swiss stock market index, and the best performing market, which I think may have been Korean 10-year bonds. Now let me overlay that with a statistical uncertainty, and actually you can now see that you cannot distinguish between these markets. The uncertainty is too large.

Rob Carver: 01:28:51 In terms of the question of saying, "Well, which market should I trade?" Actually, expected returns should not come into it at all. You should be looking at things like costs, and as a smaller retail trader, you should also be looking at the minimum capital required to trade the market properly.

Rob Carver: 01:29:09 A big theme of the book is: I have not much capital, what is the best use of that capital? Should I diversify into more markets, for example? Should I move my trading system from being a binary trading system? which is kind of the first trading system that I introduce, partly because it's intuitively easier for people to understand, and also because it's a better use of scarce capital. Should I use scarce capital, and if I have spare capital, to implement a non-binary trading system.

Rob Carver: 01:29:43 It's answering a lot of those kinds of questions.

Adam Butler: 01:29:49 Well, that is a very nichey book. Very interesting. But it's not what you're focusing your research on?

New Research Focus

Rob Carver: 01:29:56 No, my research at the moment actually is not really research, but I've got this open source version of my back-testing platform, which is called Pi System Trade. I'm currently in the process of adding to that all the bits and pieces that will move it from being a back-testing platform to a back-testing and execution platform.

Rob Carver: 01:30:23 I'm doing that for two reasons. Firstly, the platform I'm actually currently running my trading on I wrote about, well over five years ago now. It's starting to fall apart through lack of maintenance, and as a developer you often decide it's much easier to rewrite the whole system from scratch, than to try and fix up what was there before. That applies even if the person who wrote the old system was yourself five years ago.

Adam Butler: 01:30:50 Yep.

Rob Carver: 01:30:51 But that's a precondition then for doing more research. Because, at the moment, if I do come up with some initiative trading strategies, the old platform is not really sufficiently flexible to trade those. One thing I want to do is build in a lot of flexibility around things like, for example, trading with multiple instruments, so creating synthetic spreads. That's one research area I want to look at. And I also want to look at mean reverting strategies, because at the moment, I've got a big focus on carry and trend.

Rob Carver: 01:31:28 I was producing some graphs in my university course to explain the concept of uncertainty, and I noticed this really consistent effect, which is that pretty much in all futures markets you get kind of good trends, with a sort of a three to six-month lifespan. But then at the shorter end of things, markets tend to mean revert more strongly, over a period of a few days. I need to implement an execution algorithm that can capture this effect, but not get killed with costs. That's kind of the research area, where that's probably a few months off because I'm going to have to write a lot of code first.

Adam Butler: 01:32:13 Right, got it. Yeah, I think that's a pretty common research cycle. We've been going through the same thing for about the last year and a half. We launched a variety of strategies on an older trading and execution platform or infrastructure, and then got to the point where the research pipes get clogged because you can't easily deploy new strategies. You've got to revamp the execution and reconciliation architecture in order to facilitate the deployment of new strategies. That takes a little while.

Rob Carver: 01:32:49 Yeah.

Adam Butler: 01:32:50 And then we're right on the cusp of a really exciting new innovation cycle I think, but it does take a little bit of indigestion there while you get all the pieces in place to facilitate that.

Rob Carver: 01:33:02 Yeah.

Adam Butler: 01:33:02 I think that's pretty natural.

Rob Carver: 01:33:04 It is, and that's exactly the cycle as well that I saw in my ... over at AHL as well. You're definitely not alone there. I should say one thing, actually, before this call started I am writing a blog post on trend following and skew. I'm going to still do a little bit of these odds and pieces on the more intellectual stuff when I get bored of writing code.

The Skew of Trend Following Strategies

Adam Butler: 01:33:27 Is it on the skew of trend following strategies with different lookback horizons?

Rob Carver: 01:33:34 Yeah.

Adam Butler: 01:33:34 And observed at different frequencies? That sort of thing.

Rob Carver: 01:33:36 Exactly that. The theoretical result that you'd expect is that you get positive skew with trend following because it's kind of like owning a straddle.

Adam Butler: 01:33:51 Right.

Rob Carver: 01:33:51 See, it's a long volatility bet, essentially. Now, and the interesting thing is that you only get that when the time period you're measuring the skew over is tied into the time period that trend following is happening on. For example, if you are looking for trends that last for a few months, well there's absolutely no reason why it … that your daily skew to be positive, because on a daily basis the skew you're measuring is mostly the skew of the underlying assets.

Adam Butler: 01:34:24 Right.

Rob Carver: 01:34:24 It's only when you get out to maybe annual returns that you see a positive skew, unless you've got a ... but it was if you're trying to pick up really short trends, you've got just a few days, which by the way is probably not a good idea, because it doesn't tend to make money after costs. But on a theoretical basis at least, you can then on a weekly and monthly basis see the positive skew coming in.

Rob Carver: 01:34:45 Now, this is all well known, but the theory I have which I must say I haven't yet tested, so don't write this down anybody and trade it. Generally speaking, there has been literature saying that owning negatively skewed assets its profitable. That makes sense, because negative skew is something that you want to avoid. So I'm

wondering if trend following strategies have a bias towards owning negatively skewed assets, because they don't on average go up more.

Rob Carver: 01:35:19 An example would be of short volatility, which is the most negatively skewed thing you can own - so if you're short VIX futures. A trend following strategy will have a bias towards being short VIX, which means that it's daily skew will be pretty horrendous, because it will be basically picking up the skew of the underlying asset. That's the sort of relationship I'm going to explore, and hopefully find something interesting.

Adam Butler: 01:35:47 Interesting. Have you explored ... there's been some papers on the ability for skewness to act as an indicator, or as an edge, where you want to have an emphasis to be long assets with large negative skew, and short assets with large positive skew, and measured over sort of intermediate horizons. We've done some research into that that looks promising if you don't define skew exclusively on a traditional daily basis. Have you looked into that at all?

Rob Carver: 01:36:26 That's one of these things that I want to ... I've seen the papers, and I definitely thing the effect exists, and that's what inspired me into this idea that maybe one of the reasons why trend following has perhaps even negative skew at shorter horizons, is because this effect is biased towards buying assets with negative skew. That's on my to-do list.

Rob Carver: 01:36:52 One thing about this is whether, I think it may be confounded with carry in some cases, so if I think about it as a time series rather than as a cross-sectional effect. Should you load up on an asset when it's skew is more negative than normal? I think that maybe some ... maybe you guys have looked at this, but I think there may be some overlap here with value and carry, because if I think about, say, 2008, when the skew of equities and short volatility was obviously very, very high, which suggested a big long position. Both a value and a carry strategy would also have suggested a big long position there. It will be interesting to see whether the effect survives once you do your factor regression and throw those things in there.

Rob Carver: 01:37:42 But I'm with you, that with the kind of idea of adding lots of small things to the portfolio, which may have weak effects. I think it's ... to me it's a conceptually and intuitively strong enough idea that even if the effect is weak I'd still probably include it in my portfolio, absolutely.

Adam Butler: 01:38:02 Yeah, I mean our observation is that the Sharpe ratio of a diversified futures skew strategy is on the same order as what we observe for diversified carry and diversified trend.

Rob Carver: 01:38:13 Yeah.

Adam Butler: 01:38:14 And that the correlation is somewhere in the neighborhood of .3, .4-

Rob Carver:	01:38:17	Yeah.
Adam Butler:	01:38:19	... relative to both trend and carry.
Rob Carver:	01:38:22	Yeah.
Adam Butler:	01:38:24	It acts as a reasonably nice diversifier in the portfolio. Just again along this theme of adding reliable small edges, which when you apply those small edges to a sufficiently diversified basket of instruments, and in complement to one another, can actually produce a reasonably strong long-term performance profile.
Rob Carver:	01:38:47	Yeah. I just think we ... the only note of caution I'd strike is that potentially you could be ... if you've got short volatility in your portfolio, and I guess equities to a lesser degree but, by including carry in your portfolio and trend following in your portfolio that's allocating to volatility, and then including negative skew. Because these things will all ... unless you're doing it on a kind of normalized time series basis, where you're only looking at the skew, or the carry, or the momentum relative to the history of that asset, or relative to the asset class, perhaps.
Rob Carver:	01:39:26	You're probably going to increase your exposure to the short volatility factor, or the equity factor, whatever you want to call it.
Adam Butler:	01:39:34	Just make it ... the strategy, more procyclical in general.
Rob Carver:	01:39:36	Yeah, exactly. As long as you don't overdo that, obviously that's a source of return, and that's good, right? Short volatility ... these things are all, in a way, proxies for the latent factor of short volatility.
Adam Butler:	01:39:49	Mm-hmm (affirmative)
Rob Carver:	01:39:51	You're picking up short volatility in various different ways. I think as long as your risk management can understand and control for that, then it's fine, definitely.
Adam Butler:	01:40:01	Absolutely, yep. A really, really good point.
Adam Butler:	01:40:04	Well, listen Rob, I've had you now for about an hour and 40 minutes, which has just been magnificent. Thank you so much for taking the time to speak with me today, and I know that our listeners are in for a real treat.
Adam Butler:	01:40:18	Thanks again. Enjoy your evening, and let's keep in touch.
Rob Carver:	01:40:22	You too, Adam. Nice to talk to you.
Adam Butler:	01:40:25	Have a great one. Bye now.
Rob Carver:	01:40:26	Bye.