Man Group PLC

Presentation

Luke Ellis

Chief Executive Officer

See if this works, good morning everybody. So this is not a Capital Markets Day just to be 100% clear. The idea of this is to try to get people to understand the way we think about technology in an asset management business, the way we use it here. What that means for the future of the firm I think I'll leave you all to draw your own conclusions.

So this is very much about things we're actually doing and the way they operate as it goes throughout the whole firm. You'll see from a whole range of people, some of whom are used to standing up in front of people in ties, a lot of whom are not used to standing up in front of people in ties. Don't judge them on their presentation skills because a lot of them the answer is if you had a coding competition with them, it would be embarrassing.

If I had a coding competition with any of them it would be embarrassing, it'd be the real answer. That's sort of the point of it. Mostly when people talk about technology in asset management and they think about it, they talk all about alpha and whether the quants are going to take over from the discretionary people, whether humans are no longer going to be able to pick stocks, whether analysts are going to be needed. It's alright, we believe analysts are needed, we believe humans can do a lot of things.

We will talk over here about what's going on in alpha generation. Alpha generation, if you think of the basic process of investing you have the process to find some gross alpha because if we keep it in stocks land, but we're not just going to talk about stocks, is to find some stocks where you know something that isn't in the price, so it'll adjust to your price.

That can be done by humans, can be done by computers. We'll talk about as we go through the course of the day some of the places where humans have an advantage, some of the places where computers have an advantage and particularly the places where technology can help humans be smarter.

But it's really important to think about the different stages of investing. You have the gross alpha, you need some gross alpha otherwise everything else is a waste of time. Just to be clear, we're in the alpha business. Here we'll use alpha in the broad sense of added value. We don't have to get too obsessive about relative to a particular index. It's about adding value in what we do rather than passive things.

So you have the alpha generation. But really importantly once you've generated a really good idea, the amount of the alpha in that idea that you harvest, which is what you actually deliver to clients, is affected very strongly by a) how you get in and out of the trade, b) how you size the trade and how much risk you're using up and how you put that together. Portfolio construction and trading are very, very important to the amount of return that a client actually gets out of a gross alpha.

As we'll show you over the course of the day, there are a number of markets today where on trading particularly computers are way ahead of where human traders can really add value. There are areas today still for instance in - so if you take credit, if you take corporate credit the reality is today at least - well the reality at the end of last year, humans added value in the trading because finding where bonds were, finding what the prices were was like an open outcry - I mean a sort of telephone process.

That's changing with MiFID II but still today in credit humans add a lot of value in talking to the street. What you'll realise as we go through is in futures for obvious reasons or FX they really don't, that most of the execution, the vast majority is much better done by technology. Similarly, increasingly the reality is if you're going to use the risk budget that a client gives you optimally technology is necessary for portfolio construction.

I sort of laugh often when you meet a client and they say we don't want quant managers, we only want purely discretionary managers. You sort of look at them and think okay, so are you saying you want a manager who doesn't use a mobile phone, doesn't use the internet. Once they've got a position they just write it down on a piece of paper and remember what it was and try to think about the amount of risk they've got in their different positions with post it notes stuck on the wall.

Are they going to know? Of course not, right. I mean everybody does research on the internet; everybody is using spreadsheets or whatever to manage risk. We've been using technology to make humans smarter for a long time. Here at Man we have tried to embrace the idea that we want to be pushing the boundaries of what you can do. We think that it's incredible the speed with which technology has improved over the last I don't know, whatever five, 10, 15 years.

Somebody will show you a slide later and I don't know who it is but we all love the statistics about the power in the Cray Supercomputer which used to be the thing we thought was coolest compared to the iPhone in your pocket. Technology is enabling us to do things today you couldn't have imagined and in five years' time and 10 years' time much more.

We want to use that in our alpha generation, in our trade execution, in our portfolio construction, in the way we interact with clients, in the way our back-office processes work, in the way our risk management processes work. Efficiency in asset management is incredibly important. It gets harder every day to find alpha, that's because the world becomes more transparent every day. But what really, really matters is the efficiency of turning your alpha into returns for the client.

That efficiency comes in the way you trade, in the way you maximise the amount. If the client gives you a risk budget of X, if you spend your time using a lot less than the risk budget you are not efficiently using the limit the client's given you. You want to be using that risk budget on average over time fully. They gave you the budget for a reason.

It's only with technology you can really understand fully the risk you're taking. All the way through to the more you can make your settlement processes and all of the other processes efficient. You take costs out of the process and yes from a shareholder point of view we can be profitable on a lower level of fees. From a client point of view you can give them a higher proportion of the return because you're taking less out in costs and fees. That is the essence of what hopefully we'll tell you as the day goes on.

If I can work this - there we go. So you'll see we have a range of people there covering reasonably all the areas of the firm. We're going to try and cover a number of things which are very sort of practical on the alpha side. We will then, so nobody's head blows up, we'll take you for a walk downstairs. You can see the trading desk in operation which is rather amusing whenever we bring clients onto the floor.

Almost the quietest bit on the floor is the trading desk. They ask how can you be trading trillions of dollars of stuff a year in a silent trading desk. You'll see how it operates and then we'll come back to talk about some of the more philosophical things about how you make this work in an asset management business. Because the way you hire quants, the way you hire technologies, the way you give them interesting things to do, manage them and so on is a very different process than the sort of typical bank type of thing.

You can ask questions all the way through. People will answer questions if it's about what they're talking about. If people want to ask shareholder type of questions, Mark and I will deal with those at the end. So if somebody I don't know, asks Giuliana about what flows are in AHL I will stand up and say shut up. So I will try and stop any of them answering a question. They're not practised at refusing your questions but I'm reasonably good at it.

So with that I will pass onto Sandy.

Sandy Rattray

Chief Investment Officer

Thank you. Okay good morning. Okay so I'm going to give you a bit of an introduction to technology driven asset management. As Luke said we really don't view it as quant on the one side and discretionary on the other side and this sort of binary division between the two.

As you'll see there's quite a lot of human decision making in quant and there's a lot of technology in discretionary and it's just not a black and white world. The way we've organised ourselves as you'll see on the 6th floor is we have AHL on one side of our floor and GLG on the other side of the floor. There's no wall. There's a trading desk between them because that felt like the best place to put the trading desk and there's an intentional dialogue between the discretionary and the quant businesses.

Just to give you a very sort of early sort of introduction. So similar to discretionary investing, quant strategies are very diverse. People sometimes say oh it's quant. There are lots and lots of different quant strategies and people that are able to do one type of quant strategy are not necessarily suited to do another type of quant strategy. Just the same as discretionary investing, so it's a pretty heterogeneous world. Sometimes people from the outside just assume it's all one thing. It's not.

It's grown a lot recently especially in the hedge fund world. At Man Group it's grown in the long only world as well and I think we've done better than the industry in the long only world and in quant world we've probably grown faster - in the hedge fund world we've probably grown faster than the industry also. But it's clearly grown as an industry and part of that has been due to performance. But part of it has been due to what we would call a reduction in algorithm aversion.

Let me tell you a little story which might demonstrate algorithm aversion. So five years ago I went to a conference called the NMS conference. It's the major event in New York for US endowments. About 50 of the largest endowments of the US there and I came along and I said I've got a new fund, got a 10-year track record. The minimum assets that we've ever run it at is a billion dollars. We've never opened it to external investors. We're just opening it now to external investors. We've made 16% a year and we've never had a negative year and on top of it nobody else does this.

Right, it should sound interesting to you. I can tell you that not one of the endowments was remotely interested in anything I had to say apart from the fact that we sponsor the Booker Prize. Now today that fund which is called AHL Evolution is closed and our current problem is that people pester us for a capacity all the time and say we know it's closed but can you open it just for me?

That's a pretty good illustration of how the world has changed. Five years ago I just couldn't and the reason they weren't interested is because it was quant. The private equity fellow in the room who'd been down about 50% in 2008, if you actually looked through his valuations, everybody was fascinated in what he had to say. So the world has changed and actually as you'll see quite a lot of our strategies today are closed to inflows. That is as a result of a reduction and what I'd called algorithm aversion.

People's fear of algorithms has very clearly gone down. Algorithms are an embedded part of our lives and I'd love to know who got themselves to our office today for this presentation using a paper A to Z. I suspect none of them, 25 years ago that how we found how to get around town. Together we all use an algorithm called Google Maps or something similar.

The other point is that improvements in technology both hardware, software and the amount of data really facilitate this growth of quant. I'll give you some examples of that, but this is not sort of just words, it really is a very big difference in terms of how that industry has moved along.

Final point on this first page of introduction, if you want to insult us, if you'd like to insult us then the best way you can do that is to call us all a black box. So I look after both our quant businesses as well as our discretionary businesses and I must say humans and predicting what humans are going to do is often much harder than predicting what an algorithm is going to do. So we will be very transparent with you, we are very transparent with our investors about how our algorithms work.

If you need to call it a box that's up to you. But if you could call it a transparent box, rather than the black box I think that'll be helpful.

So here are some, I don't think this is actually exhaustive of the types of quant managers. There's a lot of words on this page, so I'm going to take you through them a little bit. We range from high frequency traders; those are the people of Michael Lewis' book. He would say that they're a force for the worst and I will talk a little bit about that. So these people trade equities, futures and FX. They have extremely short holding periods and they are very, very sensitive to technology infrastructure.

I have a friend who runs one of the large ones. He tells me that they have to spend $100 million a year on hardware just to keep up. They can't take a year off. That is almost unique to the other businesses. We don't do high frequency trading and our sense is that the high frequency traders really have taken over market making, that human market making in listed instruments is becoming a smaller and smaller share of volume. So the human trader of old really has started to disappear.

You may argue this is a good thing or a bad thing. Michael Lewis says it's a very bad thing. I think what I would use as a counter-example to that is if you went to the London Stock Exchange in the 1960s when it was all on a floor, there was one woman, the rest of them were men. As far as I can tell most of them went to school together and they all wore a funny sort of hat to operate on the London Stock Exchange.

Those who think that's a better world and they made quite a good living out of it, I would respectfully disagree with. I think we've got to a world where spreads have come in dramatically. There have been some learning lessons like all innovations over time. But we don't do this, we don't do the high frequency trading, holding for up to minutes. A couple of them might hold up to an hour or two.

But the next section statistical arbitrage, now statistical arbitrage is longer holding periods. It's hours to weeks typically, mostly it's in equities. It's using a series of different types of patterns, momentum patterns, reversal patterns. What you're doing there is you're now much less sensitive to superfast technology infrastructure. You don't need microwave links from London to Cornwall and then fibre optic cable under the Atlantic and then another microwave link to Chicago or something to do this. You can just use relatively standard hardware to do it.

So it's not super, super quick and the holding periods are hours. Your edge there really is finding patterns that other people don't find and doing it extremely efficiently, getting as Luke mentioned every one of the costs that are associated with this down to a minimum. We do statistical arbitrage strategies both at Man AHL and at Man Numeric. We do it in slightly different ways in those two investment engines in the firm. But that's one of our things.

The next bit is trend following. So trend following as you know was AHL's original business. Holding periods are days to months. We started off just trading futures, then we added FX and then one of our big successes is adding all sorts of new markets and we keep pushing very heavily on that and you'll hear about that later on from Giuliana today.

The core model is a momentum model. It sort of assumes that was going up will carry on going up. It sounds terribly stupid apart from the fact that it seems to work better than value over long periods of time. So not that stupid. The macro hedge fund manager's sort of maxim I suppose was the trend is your friend. That's what trend following is systemising effectively.

Again, we're looking for patterns in the data, it's not fundamental investing. It's looking for statistical patterns and that is the core business of AHL. So that's strategy 3, trend following. Number 4 fundamental macro, so fundamental macro a whole variety of people do this. Typically it's in the futures markets only. Probably the world's most successful firm in fundamental macro would be Bridgewater that has effectively built a technology structure around a set of mostly economic or financial fundamental inputs to forecasting macro markets.

So it's no longer technical, it's fundamental. The hold periods are weekly to quarterly, so a week to a quarter. So they will hold for relatively - as you can see as we're going across the page to the right, the holding periods keep on extending. It's a pretty broad set of models they use.

You're not starting to hire at the first end of the high frequency traders, you're hiring very high-end technologists that know how to make stuff work really, really fast. As you're going down to the right here, here you're hiring more macro economists who can write some code, who view that just as a way of getting things done. So you're moving less from a sort of hard core technologist and more to a fundamentals type person.

Then the last column and we do a bit of this at AHL and we've increasingly done more of it over time, we actually have some economists at AHL which is something we didn't have a few years ago. Then finally quantitative equity where obviously the markets we operate in and their equities, this for some people was the definition of what quant was. It's holding periods of months to quarters. A whole variety of different models but value and momentum and quality would probably be some of the keys ones in there.

So quality of companies' earnings, whether they're cheap or not, momentum typically would be earnings momentum rather than price momentum. Key skills would be company evaluation. So in this sort of business you find quite a lot of people with MBAs and that's pretty different to - you know you won't find anybody with an MBA in the first business. It's less technology focused. So speed now starts to become much less of an issue.

Again, the unifying thing is that people write code in this business and what we sometimes say about for example AHL is we have I think today over 30 different nationalities in AHL, about 150 / 160 people. So you've got 30 different first languages across that set of people and but one coding language and that's really important. Because if we had 30 different coding languages then nobody would understand what anybody else was doing?

We have unified ourselves in the quant businesses at Man by having everything in one coding language, something we'll talk about later on. The other surprise to you should be - so this is the core business of Numeric. AHL does a bit, but GLG I've put under there as well because what you'll see from a presentation later on this morning is just how much quant there is within the discretionary business today. We've really pushed that envelope quite hard I think.

Okay, so a couple of things about underlying themes in the world that are driving this. The first is data and the amount of data is growing. I think everybody knows that. You know Facebook seems to know it for example and they've created quite a lot of the data. There are questions about what's right to do with data and what's not right to do with data which have arisen and they affect us as well as people in social networks.

But this quote from Hal Varian who wrote A level economics textbooks a long time ago and now works at Google which tells you a bit about how the world has changed. Between the dawn of civilisation 2003 we created five Exabyte and now we're creating that amount every 2 days. So the growth of information has been massive, but the other thing which has been very important is these blue bits are structured data. Structured data means it's in nice tables.

It's nicely organised, it's easy to access, it's easy to run some sort of query on to find some sort of pattern in it. The yellowy bit is unstructured data. This is text and it's images and it's videos. So the really big growth is in data that's actually quite hard to use and I think again what you'll hear later on is how we've made advances not just in the blue bit which is sort of the bread and butter for quant but how we've made advances in the yellow bit.

Our machines read, it's generally a way that we can fairly successful upset equity analysts to point out that we have machines reading their reports. But we do, we have machines reading their reports. They don't read very well but they read a heck of a lot and they read everything. They have perfect memory and they can read them all and work out whether the analysts are becoming more positive or less positive, how they compare to other analysts.

That is a very powerful thing to be able to do and the machines will get better at reading. They're not that great today but it's a steady process of improvement as you've seen in many other areas, steady process of improvement.

The next sort of theme I think, so more data, the next theme is the cost of storing this data which is this table here, has gone down. My little anecdote here is that I spent 15 years working at Goldman on the trading floors and one day I came into work in the late '90s and my boss was very cross with me. The crime that I had committed was storing too much data and storage was expensive in those days. We seemed to pay a lot more than this \$100. So I think we paid \$1000 per gigabyte per month to rent this very high-end storage stuff. Well I hadn't stored one gigabyte, so I guess that was the source of the complaint.

But now we pay roughly a cent per gigabyte for storage and we don't bother with high end storage. We just have several copies of it so that if one lot goes down another lot is there as a back-up. The idea that I would come in and blame one of our junior analysts for storing too much data is so far-fetched I couldn't even imagine it. You might imagine being upset with an analyst for not storing enough data, like not capturing stuff and letting it get lost.

I think the processing power chart sort of showing by year the number of transistors in a CPU it's fairly obvious, this is Moore's law, I think it's pretty well known. There's a more graphical way of showing this, so this is the Cray-2 computer, it was the world's most powerful computer between 1985 and 1990. I wanted my parents to buy me one of these things. It had 2 problems, it cost \$32 million and it wouldn't fit in my bedroom.

It could process 1.9 billion floating operations per second and it had a 2-gigabyte memory. So the iPhone 7, which is nowadays not a particularly new device, has a CPU that has roughly 100 times the processing power and has roughly 100 times as much memory. So it's dramatically more powerful and the sad thing is that people use this thing to play Candy Crush and not do something more useful.

But hopefully you'll hear that we have worked out how to do more useful things with technology than some of those things which are very processing intensive. But this has been a revolution right, it's been a really significant change. Our job is to stay on the forefront of this. It's not easy. There are lots of things you try that don't work. You'll probably hear a little bit later on, Machine Learning one of our really big investments didn't have a very good February. Does that mean you give up on Machine Learning?

It's a little bit like, in my view, Uber has a self-driving car that unfortunately killed somebody in Arizona. It's big news, self-driving car kills somebody. If you think about the number of people that were killed by cars during that month, unfortunately it's a very large number, but that is news and we think about this a little bit the same. We will have some accidents along the way. Hopefully they're minor, hopefully we have good ways to learn from and deal with those accidents.

But this is going in one direction in our view. This increase in data, this increase in processing power means that we really are in a fantastic period for finding new patterns in financial markets. Our challenge on the discretionary bit is we still think that there's lots of need for human managers. You know reading is an important skill and our machines read really badly. They read a lot, but they don't read very well.

If you're a fixed income manager for example, the location of a single comma in a single clause can dramatically affect the valuation of a bond. Our machines are far, far away from being able to pick up that degree of subtlety. But they can do breadth and they can do more breadth and more quantity than humans can ever reasonably do. I think our job in the next few years will be both growing our quantitative businesses here but also increasingly equipping discretionary managers with quantitative technology to make decisions better.

It's giving them more and better information so that they're not doing manual processes. They're actually using the bits that humans do really well and delegating more and more to the machines. As Luke said an awful lot of stuff has been delegated already, this is not new. It's just a path that we're on the course of.

So that's it from me. You can ask any questions you like, so long as you don't call this a black box. So I'm very happy to take questions now or if not, I'll pass over to Shanta. Yep?

Unidentified Speaker

What do humans do better than machines?

Sandy Rattray

What do humans do better than machines? Well, I mean a whole variety of things quite clearly today. So there's the bond where a particular clause make a lot of difference to whether the bond is worth a lot of money or worth nothing and machines are very far away from that. Meeting the CEO of a company and working out whether the body language says the growth plans are pie in the sky or something real, again machines are pretty far away from doing that.

I'm a little sceptical about forecasting I must say. A common criticism is well the machines just look for patterns in the past and the humans forecast. Well humans do not have crystal balls. By and large they're using data from the past to try and intuit something about the future. So I'm not sure that prediction is something that either machines or humans do very well.

So it's often the softer part, either dealing with people or sort of understanding very specific features of a company that humans do better. In our experience we found it very hard to get humans to do a good job of forecasting market direction. That's one of the things that we just can't seem to make any money out of. So machines seem to actually be better than that. The signals are very weak and you need a lot of breadth.

On the other hand, deep dives into individual companies, that's the preserve of humans today. So that deep dive but then we found that once they've dug deeply into a company machines can help with an awful lot of other aspects of it. Position sizing, risk management, execution we found were all things that are very well done by machines.

So it is genuinely a complimentary world right now. But sometimes if you look for example at Black Cab drivers are blocking parliament square in protest against Uber, I don't think that's a good response to technology. I think if I was a Black Cab driving I would be trying to work out some ways to use technology to do a better job or employ more people or do something else. But I wouldn't be trying to protest against it. I think it's coming and the question is how you navigate it and use it as well as possible.

Unidentified Speaker

Hi, your comments on the rate of change in technology were extremely interesting. Luke said that he felt that over time as transparency increases the alpha opportunity in markets is reducing if I've interpreted that correctly. So where does that end up if over time transparency increases and alpha declines?

Sandy Rattray

Okay that's a great question. I think two things, number one I agree with the comment right that alpha declines as transparency increases and on top of it, you know damn it the hedge fund industry keeps on growing. So the number of dollars chasing that alpha has gone up a lot despite what the newspapers say, we're at peak assets and we haven't shrunk for a decade.

The first thing is do what you do more efficiently and I had a funny conversation with the head relationship manager at one of the banks that covered us who wanted to know what his bank's relationship with us would look like in the future. What he actually meant was who do I talk to for trading activity in the future and I said well the person you talk to is a box because the box is going to work out whether Morgan Stanley or Deutsche Bank or JP Morgan or Merrill is giving us the best execution.

There's too much data for humans to work it out. So efficiency is incredibly important in particular in the hedge funds world. I don't think it's a very efficient world. I think it's been historically a world that said we've got fantastic alpha and so costs that's for other people. I think that moment is far behind us, costs is for hedge funds just as much as it is for anybody else. So that's one thing.

Then I think the other thing is that there actually are a lot of new sources of alpha out there that have become available to the people that get there first. I view this as a world where what was alpha yesterday has largely become beta today and therefore you need to keep on finding alpha because it'll start to get replaced by other people working out how to do it.

To give you a feel the AHL momentum strategy started in 1992 and the fee on those strategies was 10 and 20 in 1992. Not 10 basis points, 10 points. So 10 and 20 on a couple of hundred million dollars' worth of assets right. Obviously it's all been downhill since then, but now you get to charge much, much less on relatively simple momentum models. They're in textbooks. So it's become closer to beta than alpha.

But there are lots of places and you'll see hopefully over the course of today new places where we are finding alphas and those will get eroded in five or 10 years as well. So it's a continual battle to keep finding them. But I'm not pessimistic. I actually think this growth of data thing gives us enormous opportunities that we didn't have in the past. Okay that's me. Yep okay Shanta.

Shanta Puchtler

CEO, Man Numeric Boston

Thank you, Sandy. So, my name is Shanta Puchtler. I'm the CEO of Man Numeric Boston and I think I'm the only American up here today. So I feel a little bit in the minority.

So what is Man Numeric? Simplistically on the left side of this page, we take interesting data; we build fundamental hypotheses off of that data about how we can invest in equities. If we like what we see and we think it's a valuable signal in our process, we implement it. So that is all about modelling, portfolio construction and trading, all of these components that involve technology, involve sort of efficient execution of that risk budget in the overall investment process.

The right-hand side here is a typical buy candidate at Man Numeric and simplistically we like cheap stocks that have positive characteristics. So we sound very much like a discretionary shop when we say that. The devil is in the details obviously of the models that support that. So we look for attractive valuations, favourable momentum on a whole variety of metrics. We look for strong management and we measure that in a variety of different ways associated with the company. Obviously not by interviewing the CEOs but by drawing out a whole bunch of quantitative data about the behaviour of that management and their execution style.

We look at how other investors treat the stock in the market. So we look at the options markets, we look at short interest; we look at these other sources of information associated with the positions that we're interested in taking. That's a high-level view.

Giuliana Bordigoni

Head of Alternative Markets at Man AHL

We changed topic, so what I would like to talk to you today is the process in which we evaluate and access new markets. The idea, what I hope at the end of the presentation is that you get an idea of how important is technology for us in order to be able to push the boundaries and really add many drivers to our portfolios.

So I would like to start telling you a little bit the history of AHL. You might know the beginning, so it all started in 1987 and they started trading futures. After a couple of years they added the FX market. But it's only in 2005 that the push to a new market to become stronger with the creation of the AHL Evolution program. The idea was simple. It was let's use our core models and let's apply to different markets so that through diversification we can improve the return of our portfolios.

In 2005 we started trading - well I wasn't there - but we started trading CDS indices, cash and cash bond and through the years we have added many asset classes. We can think of cash equities, interest rate swaps, mortgage backed securities, OTC commodities, I mean electricity, like German electricity or French electricity.

The latest addition is actually in January 2018 which is crypto currency in the shape of a Bitcoin future as you know is listed on CME. But this is not the end of the timeline. I do hope that if I give you this presentation again next year you will see that in 2018 we will be adding a few more asset classes, but you have to wait to know which ones.

So quite extensive experience, I mean more than 10 years of trading fixed income derivatives. About 10 years' experience in trade OTC commodities. And what does this bring to us? So what we want to do is apply our core models and leverage on our technology and infrastructure.

How do we start? It all starts with an idea. It can come from anyone. It can come from myself, it can come from my team, but it can come from anyone inside AHL, from the traders, from the operational team. How do we make sure that we actively are engaged - we're finding new ideas?

The way in which we do it is that we have like meetings with industry, we have the meetings with the brokers and we know what's going on there. We know where the liquidity is. So once we have one idea we look if it's viable. What does it mean, viable? Viable means is it enough for us to trade it daily? At the moment all our funds needs to be trading at least once a day.

So we need the liquidity that allows that. Then it needs to be diversifying. We are not yet trying to find - we don't want to add a new market for the sake of it. What we want to do is to find new drivers, which means that if we are just finding a way of accessing the same drivers that we access as well but with a wider bit of a spread, that's not for us. So that is automatically discarded.

Then of course, can we trade it? Does the legal environment allow us to trade it? What are the tax implications? And as soon as you move out of the traditional space that becomes more important.

Once we decide if a market is viable, then a more thorough investigation starts. Of course, it would be like a prioritisation and so on. What do I mean by a more thorough investigation? How do we execute this market? How do we price it? Can we price it? Can we apply our core models? So now our philosophy in general is that momentum exists in any market at different times, and the best that we can offer, as I said already many times, is diversification.

Now the point is that, there are examples in which I believe that momentum is not really applicable, and these are markets in which there are loads of interference and therefore a momentum system can't really have an edge. Then, of course, it's going to be a matter of like the legal team to allow us to have all the documentation in place.

So once it's all done, it's time to go and implement it, meaning it needs to be set up in our infrastructure. We need to get - we need to decide in which fund it's going to go and even the choice of the fund is fairly systematic. So there will be a few rules and the rules will determine in which fund it will go. Then, of course, we'll have to give it an allocation. Once it's all done we test trade it with internal capital. This is just to make sure that all the pipelines are in place.

We don't want to do this - we don't want to risk anything with client money. So once we have test traded it with our internal capital we are ready to go and we are ready to trade it, in one of the funds.

Here I can conclude; I hope that I gave you an idea of the technology that we have at AHL to add new markets and new asset classes. It's a tough game, especially going forward because, of course, the easier wins are already in the portfolio, so it's a matter of really keeping momentum in the sense of trying to find new asset classes and new additions and each of them will require to develop even more our technology and infrastructure. We are committed to do all this work and we do have an infrastructure that supports us. I'm happy to take any questions.

Luke Ellis

Thank you, Giuliana. Brilliant. So Fabian, come - so now we're going to talk onto the discretionary side, but actually particularly about portfolio construction using technology, if you can make the technology work. Just while Fabian gets up, I mean there's a fundamental belief in AHL that all markets have trending properties. So we don't take markets out when they go through a period of not trending, because actually in the end all markets have trending properties. So it's profitable when nobody else is doing it, that's definitely a benefit but it's - you know you can pick, I don't know whatever dollar, sterling, which has been around for hundreds of years, it's very profitable to trade over time. It's just you get periods where it works and periods when it doesn't.

Fabian Blohm

Co-Head of GLG European Long-Short

Good morning, my name is Fabian Blohm, I'm one of the co-heads of ELS which is to Luke's point here is, traditionally a very much a fundamentally managed, bottoms up, multimanager platform. We currently today have 20 plus PMs that we give risk capital to and we feel that our edge is in picking stocks. But on the topic, we also know full well that the individuals on the platforms are humans that comes with the flaws that humans have and we are and have over the last two to three years, invested heavily in technology to sort of create these synergies of that stock picking process using technology which have come from AHL and Numeric.

One of the reasons why we want to use technology is that technology have no emotions and at times people can make poor decisions based on emotions. PMs can at times be biased. Computers are less biased. The other problems that the humans have is solving very complex problems, in particular on topics of portfolio construction which is an effort that we have focused heavily on over the last two years. So again the computers come to our help and assistance in building better portfolios. So what we are trying to do here is use the human to source the alpha because we strongly believe that the man has the edge of picking stock for now, but the computer is much better at extracting that alpha in an optimal way. In a very capital efficient way.

So if we look at the three processes where we heavily now invest in technology and where we use technology is on portfolio construction, execution and centre book. If we start with the portfolio construction and risk management aspect of this, and we focus on this chart to begin with - it might need a little bit of an explanation here. But basically what the chart shows is the decompositions of the returns that are being generated by our UCITS fund. What stands out is this yellow piece which is basically the return coming from the fundamental stock picking component from the PMs.

The colourful lines below or around zero are systemic risks that we took to extract that alpha. So if we buy Vodafone, unfortunately we need to buy UK, we need to buy Telecom, we might need to buy value, we need to definitely buy size. But we buy Vodafone, not for those four reasons, we buy Vodafone because we think Vodafone, relatively speaking is going to outperform. Now what do we do with the other four components if they add no value, they also potentially are making the capital allocation inefficient because if they generate no value, they should get the least amount of risk that we can throw at it, to make the process very capital efficient.

So where portfolio construction and technology comes in, is to make sure, and again this chart might need some description as well, is now the risk decomposition of those decisions. So in the past, when we didn't use technology to the degree that we would like to, we allocated too little risk to what really was the edge of the business. So where we now use technology very effectively is to making sure our vol budget is aligned with our edge budget.

Technology is very critical in that ability to extract that alpha in a more optimal way. So again come back to the sourcing of the alpha, it's fundamental. The extraction component is heavily reliant on technology. Now there are other benefits here from technology, one is execution, the other one is centre book. If we continue on those topics a little bit after this actually.

So the way we more specifically extract in a capital efficient way, we deploy technology through our discretionary, proprietary system called Dragon, so a screenshot here, that we lever, using again AHL technology, to the fingertips of the PMs but also to the co-heads myself and Neil, so we can risk manage the fund, but also the PMs can risk manage their portfolios and their risk capital effectively. So this is sort of a component of us working with the PMs and helping them in a very tangible way and guide their process of how they can squeeze out incremental systemic risk.

So take the Vodafone example, they might want to buy Vodafone, but it might give a UK andTelecom exposure. Okay what could be the perfect substitute? Instead of taking down the alpha position, maybe that could be to short BT. So now you have a fairly good bet on because you have very little UK exposure, you have very little Telecom exposure, you have very little value potential exposure and very little size exposure, but you are now monetising the spread. You're monetising the alpha. You're trying to monetise the alpha in an effective way.

So here, we use technology to guide and use what we call marginal contribution that gives us this edge that we can say, this and this and these trades, gives you the biggest bang for the buck, i.e. with a small marginal change to your alpha, you can reduce the systemic risk materially, risks that haven't generated return, long term.

The other thing is that at the top level, we also have the ability to supress and reduce our factor tilt in the book. We heard the presentation previously talking about momentum. A common bias among our fundamental PM is to buy stocks that have gone up, or short stocks that have gone down.

Now that is not what we're trying to accomplish here. The aim of the game is to deliver alpha. Momentum is a by-product of those decisions. We want to supress the momentum risk in the book to be as small as possible. Partly because it's not a capital efficient way of doing it. It's better to hand that money over to my previous presenter, because that's a much more effective way of investing in momentum.

So we want to reduce the momentum exposure at the top level, assuming that's a common bias among PMs on a platform. We hedge that out and allocate the risk back to the individual book owners that carry that type of risk.

The third component here is what's called centre book and I'm actually going to flip here a little bit. The centre book is probably - the strategy that probably epitomised more or less, the optimal use of the skillset of Man Group. Centre book is a systemic trading style, just sitting within ELS. So what it does, it utilises the discretionary alpha, but what it also does, given that it's a completely systemic trading strategy, it uses the technology from AHL to implement that strategy and to link it also back to another part of the firm. The first signal of this systemic strategy was actually created by our colleagues at Numeric, because this is a signal based systemic strategy that uses the discretion decision by the PM to optimally extract the alpha from the PMs in a very capital and cost-efficient way.

Capital in the sense that, this book in itself carries around 85% of its risk aligned with the alpha generating component, and it's doing it in a very cost-efficient way because it's implemented in a very passive way.

So this is sort of a perfect example of how GLG and centre book in this case is really effectively levering the core skill of the whole Man platform.

Now I'm going to go back one page here, of course, execution is also a central component to the improved profitability of ELS. ELS is a high turnover strategy. We trade a lot. So that's why it's important that every time we interact with the market, we do it in a smart way. We don't scream and shout that we want capacity because in that case momentum or the market impact is going to be materially higher. So we need to be very smart of how we execute. With the centralisation and execution, which I'm going to talk less of because that's not my expertise, but as a consequence and a by-product of that, we are now sitting down on a quarterly basis with a TCA analyst that evaluates all the execution decisions by individual PMs. Sitting down with the PMs and their corresponding execution trader, to provide very clear guidance of how to improve execution.

I'll just give you a very clear example, you have the PM, they are doing this fundamental work, he looks at the share price, the share price is moving away from him and instead of thinking through what should I do, they go and ask the execution trader to fill the trade ASAP.

Now they might want to trade 25% of the liquidity in a very short period, which obviously materially you want to push the stock further away from them and increasing our own market impact and reducing the profitability of all our trades.

So in terms of the execution that also helped us significantly to reduce the cost base to implement the strategy. I think, I'm not wrong, in saying that the costs were actually reduced by close to 40%. This is sort of an expensive strategy of trade, so obviously this is an important improvement.

So the process of embedding quant and technology in a discretionary strategy paid off materially. And with that, I open up for questions.

Unidentified Speaker

So in your example, Vodafone and BT. You were - I understand you wanted to concentrate your specific risk here, you were taking out UK risk and other stuff. But you were at the same time taking on BT specific risk on the other side. So how would you - how would you deal with that?

Fabian Blohm

That obviously needs to be in the sense that you don't short in BT unless you've done your fundamental work on BT clearly. So that goes without saying. But you could have hedged with an index. The problem is that that's not capital efficient because basically you generate no alpha in trading in index against it. You might actually introduce other risk that you have less control over. So it doesn't mean that we all pairing trades up, that's the wrong message, but what we're trying to do is earn the spread.

Unidentified Speaker

Yep, thanks.

Unidentified Speaker

Thanks, it's a related question. I just wanted to ask about PM incentivisation in the piece, particularly with Dragon, how do you go about that? Are you incentivising PMs behaviourally to increase their stock specific risk component or are you still paying them on the overall fund effectively and allowing them to take the types of risks, but just allowing that to be more transparent?

Luke Ellis

We don't want to get into specifics of exactly how we compensate people for obvious reasons, given everybody in the room but the goal is to incentivise them to take more specific risk, and that's what they get compensated on..

Fabian Blohm

Also to answer that question, without disclosing how much we pay them which is less relevant for this. This chart is not dissimilar for most PMs. We allocate on Sharpe to the PMs. That means that we combine the risk and return from this, as well as this. So all else equals, the more specific risk you take, the higher the Sharpe and the bigger the risk allocation. So in principle terms the whole set up is to incentivise people to do the right decisions.

In addition, I mentioned here, that the factor hedge is a hedge that we do at fund level, but we allocated back to book level. We simply say that the person, all else equals, they have more specific risk, get less allocation. So again, we're trying to sort of build in the incentives of saying, you do the right thing for the fund, and for yourself then you will get more risk allocation.

Luke Ellis

Thank you Fabian. We're now going to do something slightly risky, which is to move you all. So you have to wake up to try and - make sure everybody's awake. So we're going to go down to the sixth floor. Please don't try and lose whoever's showing you the way to go, otherwise it will get chaotic. But we just want to go and show you the trading desk and a demonstration that we give to clients who obviously we have in a lot more often, who are asking about what does all this mean in terms of execution. So we have a little demo on a - it's just much easier to do downstairs than up here.

Luke Ellis

Chief Executive Officer

Thank you again. If it wasn't clear, that was all live trading. Because it's live trading, it's a bit hard to pick examples, because sometimes they work and sometimes they don't work particularly well, because it's live execution going on. The next couple of things talk about one of the key elements of that is, how do you get the really good people to want to work here. Because all that clever stuff only works if you've got really good people. Who's next? Antoine.

Antoine Forterre

Co-CEO, Man AHL

Good morning everyone. I'll try and avoid it being too fluffy, because if you talk about people you might raise your eyebrows and say, oh, another talk about culture. If you really look at what we're doing, we do two things. First, we run a number of proprietary softwares, which we develop, and based on a collection of data we push trades and execute those trades. That's the first part. The second is we continuously develop new trading strategies and also improve the ones that we trade. Those two things are really the systematic aspect. That's the first one. And the quant aspect is the second one. They don't necessarily have to go together. You can think of businesses, like ETFs, for instance, that are fully systematic, but very simple quant products.

Conversely, if some of you are in banks, serving in banks, you have strats, or quants, in banks that use all the quant techniques that we do but ultimately, they push trades discretionarily. Specifically, what we do is - it goes, hand in hand, together. They don't have to go together, but also they sometimes conflict with each other. Because what you need to run a systematic system is very different from what the researchers need to research on your system. Some of the key aspects are here on the slide. The key point of the next 10 minutes that I want you to leave with is, in order to make this work you have to not only have the great data that Shanta has talked about, the market access - and we'll see some examples, the hardware and software that Gary will talk.

This shows you how we actually recruit on board and manage people. What you see at the top, the key phases, would be similar to anyone in a bank or a large institution. That's not really what differentiates us. What differentiates us is how we go about this, the initiatives that you see below, and the level of focus that we have on the team. We, a few years ago, built in-house, within AHL, both the recruitment and talent function. They now report directly to the businesses and they're very much aligned with what we do. I'll go into a bit more detail on the recruitment and culture in the last few slides. Just to give you an idea of things that will differentiate us, maybe, from large banks, or large asset managers, just this weekend we had over 100 people in this building doing a hackathon, basically, developing the open source Python ecosystem.

That's something that most banks, probably, would not run. It's normally run by tech firms. Another example - we host meet ups. Meet ups are informal gathering of people interested in a particular topic in technology. We host two - one on machine learning, the other one on Python. Those are monthly events, where we get between 150 and 200 people to come here. They don't listen to us. They listen to anyone that has something interesting to say in the space. It's great for us, because we get the connection with 150, 200 who are people interested in what we do. We remain at the forefront of what's happening in those fields and then, eventually, that builds the pipeline to recruit. All those things are fairly different from what you'd expect from a - maybe a traditional asset manager. They feed into the culture that we try and bring.

Looking at bit more in detail at what we do on the recruitment side, those are some of the initiatives that we worked on with the broader team over the last few years, to try and build a long-term pipeline and a short-term funnel of candidates. Just a few that I want to highlight. A couple of weeks ago there was the European Maths Olympiad, girls' maths Olympiad, in Italy. We've been sponsoring, for the last four years now, the UK team. That's great, because really (A) we build long-term connection with very, very smart kids, so that's pre-university. And eventually some of them will come and do internships with us. And eventually we hope that we hire them. In particular, it's a girls' maths Olympiad, which means, if you look at the diversity, one clear lack of diversity in businesses like ours is the female ratio. .

We're pleased to say that they came third. One of their candidates was joint first. Another thing that Anthony will talk about, we have a partnership with Oxford University specifically on machine learning. This is great for a number of reasons. A few years ago - a couple of years ago we did a smaller partnership with Durham that just started a course in computational heavy astrophysics. That's a way to get long-term candidates. Then another example is the hackathon that I talked about. All this translates into a fairly active recruitment pipeline. I was looking at the stats, actually, to present to Luke a few weeks ago. In Q1 we saw over 1,100 CVs.

We interviewed 230-odd of these. We ended up making, I think, 25 offers. I should say that Matthew who is a co-CEO of AHL, and myself, we make a point to meet everyone that reaches final stage, because we think it's important that they see us and understand the culture and vice versa, I want to make sure we vet them. It's a real management focus. My last slide is on the culture of the place. You've seen the floor, so you've seen the big screen. You might have seen some of the features. If you hire a bunch of very smart, diverse people, which is what I try and do, the worst thing would be to be very prescriptive. Really, what we try and do is allow individuals to define the culture and the environment that we work in, both from a work and business point of view so those initiatives here ultimately all directly benefit the business from open source.

We talked about the meet ups. We see Anthony here at the conference. Hackathons, a number of talks unplugged, to the more informal ones. An example is, every

Wednesday afternoon I have a piano lesson in the office. And along with 28 people on the floor.

It's because we realise that if you hire people that are fairly maths savvy, there's a good correlation with music. So we have a lot of musicians and there is a lot of demand for music in the office. Coffee - to go back to another thing that I mentioned before, is another example. People love coffee. We do coffee tasting very frequently. People get to define the culture of the place. It's very important to give them this freedom, because eventually that's what translates into the creativity that they're bringing into the strategies that we run. And ensure that we don't become complacent and just become constrained by the environment that we've built. As one of the key risks of systematic strategies, is you end up limiting yourself to what you know in the environment you operate in. This is what we try and avoid.

With that, I'll shut up. Do we do questions?

Unidentified Speaker

Thanks very much, Antoine. You mentioned at the top that you had, I think, it was 53 data sources, or something like that. You must have a very large number in the pipeline. How do you go about that filtering - identification of value and filtering down?

Antoine Forterre

The workflow or work process for a data source is one that we spent a lot of time over the last 18 months, because you have an increasing supply of data sources. If I'd done the same talk, probably, five years ago - I don't know the number for sure, but we would have looked, I don't know, at maybe a handful or two handful of data sources in a quarter. There's a huge increase in supply. They're all fairly complex and take a lot of time to actually onboard and process. You want to make sure that there's a clear business rationale early on, so the researchers are involved from the outset. Then we try and make sure that we get the free sample data on board as fast as possible. Usually, most data centres will give you a subset of the data for free, so you can run your strategies on, or test them. That usually was a pinch point.

What we did is, we actually created a team internally that's tasked with onboarding, and delivery to research, those data sources. Now our researchers work hand in hand not only with the engineers that developed the models, but also with the data scientists that help them onboard the data and keep the machine turning very quickly. Then, once the research is done and conclusive, it gets passed to data delivery, so people then productionise this data source. From a process that could have taken months in the past, our aim is to have this in a couple of weeks to a month, a function of the complexity of the data source. If I may, it's very important, because more and more - you have less and less large areas of capacity that are open that we can tap. There's a few that we're working on. Beyond that, a lot of the growth and the increase in Sharpe ratio will come from adding a number of data sources fast, so agility is very important.

Luke Ellis

Chief Executive Officer

Gary was the AHL CTO and is gradually becoming the man that is in charge of all of the alpha-generating technology.

Gary Collier

CTO, Man AHL

Thanks for the intro, Luke. A little bit about me. I've been with the firm for about 16 years now. My training, my degree is in physics, theoretical physics, like a number of people here. I've been a programmer, I've been a hands-on developer since the age of about 10, 11, and was part of that first wave of people who got home computers, so I've been hands-on, really, ever since then. I'm going to talk to you about what makes a good quant infrastructure. I'm going to start by asking you to reflect on this particular statement here, which you might have seen in this, or maybe a slightly different guise in the past. That's for a technology-dependent business to innovate and succeed, you either have a great platform, or you become one. Hence, the Sigma sign at the top.

Of course, we can talk about infrastructure. It's really the way you pull your different and disparate strands of infrastructure together to form a coherent whole that really allows you to deliver. What is a great platform? A great platform should allow you to pull together your core expertise and functionality and present it back to people as a coherent whole, but also decentralise the ability for people to innovate using that functionality. This could apply to everything across Man Group, because the whole of Man Group's got a great platform because we've got great sales. We've got a great ability to structure our products. I'm going to focus on the technical aspects here.

These are the three lenses I want to look at this through. If Antoine hadn't talked about people, people would have been another cog on here, because it's incredibly important to have the right people working on these three aspects of what we do. Our ability to deliver a great platform, to deliver a great infrastructure is really governed by our capability in these three areas. I'm going to talk a little bit about hardware. This isn't just tin but our compute in general. Software, and particularly what we've been doing to both build upon, and contribute, back to open source, and then how we join all that together to be agile and deliver a great process that can meet the needs, both of research innovation, but also meet the needs of productionising the ideas that we've got. Because, often, the features of a platform that can support one, don't necessarily automatically support the other.

Apologies - it's a super-complex diagram, lots of boxes. The key thing to focus on here isn't so much the contents of each individual box. What I want to get across really is the enormous amount of different pieces of technology, different pieces of infrastructure which have got to work in unison in order to provide the capability for our quants at the top, up there. If we look through the different layers in turn, at the bottom we've got our compute infrastructure, our hardware, our storage, our networking, our GPUs for the particular sorts of problems that we want to apply that technology. Then the layer on top, where we harness that tech, in simple ways, maybe just giving research or a particular headnote, or environment on which to perform interactive research. Perhaps more complex ways of orchestrating it to run large-scale distributed simulations.

The development tools that we give to everybody in AHL. Out of 170, or so, people in AHL, well over 100 can, and do, write code on a daily basis. Then we move up through different layers of the software stack, the open source ecosystem that Antoine talked to you about, that we great value from. Finally, we've got the core libraries, the core frameworks that we build on top, our IP, that harness all of the layers of the platform underneath. I'm going to start looking at the hardware side of things a little bit. As I said, it's more than just the tin. It's also storage. Key drivers for us here are performance and agility. Performance - we want to be able to do things really quickly, whether that's pull a whole mass of data back for the compute that needs to run it, but we also need to be agile. Agile here means putting the compute into the hands of the right people to do the right task when they need it.

Two parts to our strategy here. Cloud and, largely, on-premise cloud. Cloud - it's almost a given nowadays. Everyone talks about cloud, everyone talks about third party cloud and, yeah, it can save costs et cetera. The key driver for us here is the agility that cloud can bring to us. That agility can manifest itself in a bunch of different ways. It can be something simple. A new researcher has joined the team. We want them to be up and running really, really quickly. At the click of a button there's a new research environment for them that's been spun up. A lot more than that. We need to be able to use the resources we've got flexibly. Some of the frameworks which were on the previous slide, frameworks like slurm and spark, use the cloud infrastructure to run big distributed compute jobs, simulations, risk, that kind of thing. We also go quite a lot further.

We're heavily into automation of all of our compute resources. Pretty much everything we've got in terms of compute and infrastructure is scripted. We could rebuild our compute cluster from the click of a button, if we need to re-provision the entire cluster, say, onto a new set of hardware. We can run tests in a massively automated fashion. Want to run, perhaps, a complex test, say, against our trading infrastructure, which is a large suite of boxes. We can spin up a complete clone of that in our internal cloud, run a test. It could be a new piece of software. It could be an operational test to check some change is good before we release it. And then tear the whole thing down again in a matter of moments, and start again, and do something fresh.

You might ask, why are we pursuing an internal hire strategy, an on-premise strategy, when AWS and Azure, they provide clouds, right? The answer here is we want to give best-of-breed capability to the department. Best of breed in terms of the fastest flash storage that we can get our hands on, the best compute, the best networking. Right now, those needs are serviced by building this capability out internally. Things might change in the future. What I've got my eye on in particular is the convergence of data with external cloud. Most of the big external cloud providers, they are looking to - the likes of Thomson Reuters, Bloomberg and bunches of other vendors. If they start to become more of a one-stop shop that provide not just your compute infrastructure, but also the software that we need to run it, and the data, then I think that might be a tipping point at some point, in the next three to five years, to look more closely at off-premise cloud.

Onto the software ecosystem. The software stack that we've got at AHL and, increasingly, at Man Group, there are three key elements to this - Python - that's a programming language that we use; open source software; and about 2 million lines of our own proprietary code. Python, as a programming language, has been a really big success story here. We first started using it in earnest around about 2010, 2011, when we decided to completely rebuild our research and trading infrastructure and standardise on Python as a single language. Before then, we'd had different languages used for research and different languages used for production trading. This language barrier always put a translation wall between quant and between technology. Quant would come up with a model and, then, once they'd satisfied themselves it was a good model it'd be thrown across the wall and technology would implement it.

Now, everybody uses Python. It's brought our quants and technologists much closer together, so much so that you could walk across the floor downstairs, and I could say, is that a quant, is that a tech. Basically, it would be a toss of the coin whether you would get it right or not. The language is great. Python is a fantastic language for data science. It's probably the de facto language now for doing data science. Part of the reason for that is the massive ecosystem growth that was seen over the past or six years in all of the surrounding libraries and frameworks that make that possible. Libraries like pandas, for time series manipulation, psychic learn for machine learning, libraries like dask and spark that allow us to run big distributed compute jobs, there's a huge ecosystem there that makes the sort of work that we do possible.

Now, a lot of companies are quite keen to take, a little less keen to give back. We find that by actively engaging with the community, by contributing as well as taking, we can attract and retain the brightest and best developers. If you were to look on Man AHL's GitHub externally, you'd see some of the packages that we've actually promoted and are now part of an accepted data science ecosystem. The final piece of that stack, and not to belittle it, is the 2 million lines of proprietary code that we've got in-house. Really, a lot of the innovation, a lot of strength and power in technology nowadays comes from how you take accepted pieces of tech at libraries, frameworks et cetera, wire them up in innovative ways using clever people, and then feeding them with your own data. The concept of re-use in software really isn't a new one.

Twenty years ago, people used to talk about re-use and re-use classes and blah, blah, blah, et cetera. In a sense, we've reached this golden age nowadays where it is actually happening. Really, it's these things all joining up that gives us the platform power, the agility, the ability to innovate and take existing things, plug them together in new ways. Which would all be fine, but, at the end of the day, we've got to take what we build and ship it to a production - our production system. We've got to have it trading in a reliable manner. As I said earlier, the requirements for a flexible, innovative research environment can sometimes be perceived as quite different from a robust production trading system, trading billions of dollars, carefully locked down, carefully change controlled et cetera.

All of those latter things are true, but by having a very robust process, peer review of everything, and by automating as much as possible, we can bring those same levels of agility and ability to change to our production platform, as we can to our research environment. Over the course of a single month we're typically making 200, 250 planned changes to our production system with very high reliability. They could be simple upgrades of trading system components. They could be new predictors that have come out of the research team, things of that nature.

The final slide, the final thing I wanted to say before opening up for questions is, this platform, the things I've talked about, grew from within AHL. We've been heavily investing in this platform for a number of years. Over the past couple of years, the focus, well, it's not so much shifting, but the focus is to take what we've got. Take the great elements of the platform that's given us a Python infrastructure. They've given us great development environments, the tools, common analytical functions, market data access, and bring those out into a common set of foundation infrastructure to support the quant efforts across all of the different elements of the firm. So, to close by asking you to reflect on the statement on the first slide. A technology-dependent business to innovate and succeed. You either have a great platform or you become one. I think we have, and are continuing to evolve, within the firm, a great platform for doing quant. Thank you very much.

Luke Ellis

Anyone dare a question to Gary? I did say towards the end of the session we'd get to people that could leave us all for dead. Here's one.

Unidentified Speaker

Hi. Just in terms of new platforms, I think you said there's more one-stop solutions potentially being provided by Reuters, Bloomberg et cetera. Does that reduce the barriers to entry for other entrants in the market?

Gary Collier

I guess, conceivably, it could do a little. I would roll back to the first or second slide, where I showed all of the different pieces that need to come together. I talked about data a little bit, because, as Antoine said, one of the key focuses we have is our ability to move forward, considering more and more varied data sources. Bloomberg and Reuters, and other providers, perhaps coming onto Amazon are particularly of interest to us, because it's a final piece that would help us become even more innovative. You still need all of the other parts of the stack as well.

Unidentified Speaker

It was just a slightly related question there. The advantage of bringing - working with people like AWS is that - and having data sources already on board, is just that proximity in the race, the, essentially, physical proximity of data to the problems you're trying to solve and the speed, therefore, with which you can get hold of data to test things?

Gary Collier

Physical proximity is a little part of it, but I don't think it's the main part of it. What is a constant in this job is the amount of time spent dealing with what is a dirty world. When you're dealing with data, you're dealing with a really dirty world - things that are incomplete; symbology mappings which generally need to be done again, and again, and again; questions about coverage; questions about point in timeliness of the data. I think it's those things that would really help, less than just physical proximity, which is also kind of important, particularly with large datasets.

Unidentified Speaker

Did you say, by having a third party own that, that's sort of thing, it's about data protection?

Gary Collier

Yes. If one of the big cloud providers - you could say, well, it's not their core business, but Google and Amazon seem to be doing a lot of things now which aren't their core business. If someone answered that problem with regard to data, that would be an interesting proposition for us.

Luke Ellis

The trick is you're trying to get robustness and flexibility at the same time, which is hard to do. And the further away you are from the problem, you end up giving up on one or the other. When you look at, typically, whether it's a Bloomberg solution, or, in a different context, an Aladdin solution, they give you incredible robustness, but less flexibility. Actually, if you want to innovate, it's useless, but it gives you - gives the Board comfort. If you go into something which is purely about innovation, it's really cool, but you don't have the robustness to deal with the risk in there. What we're trying to get is the balance. Cool. Thank you, Gary.

Luke Ellis

Chief Executive Officer

We saved this for last, in part, because we wanted you all to be warmed up by the course of it, and in part because everybody wants to talk about machine learning. We thought, by putting it last, we'd have a roomful, which we have. Anthony has the glorious title of Chief Scientist at AHL, which I think speaks for the fact that he's now the cleverest person who's walked in the room today. Over to you.

Anthony Ledford

Chief Scientist, Man AHL

Thank you very much for that somewhat flattering introduction. Welcome. Very nice to see you all here. Just a little bit of a confession. I'm going to talk to you about machine learning, but I'm not a machine learning evangelist. In fact, when we started doing machine learning, back in 2012, I was one of the major sceptics. I could see there was a lot of activity going on in the outside world, not much in the realm of finance. It was our duty to really understand whether this was all hype, or, if it's not all hype, how much of it is useful to us, so we embarked on a study of doing that, but very much from the point of view that we were all sceptics. Anyway, we built some prototype systems, evaluated their performance, and we saw that a subset of those had some interesting features.

Our first reaction is, we've done something wrong here. We've made a mistake. We're forward-running the data in some way. We're looking at forward information. We maintained that very sceptical view. But we checked, and we double checked. We were unable to actually find any errors in this. So, much against our preferences, we pushed it into the next phase of our research process. This is where we actually build the trading system and we go ahead and live trade it in the markets, but we're not using any client money, completely safe from the client book. This is what we call test trading. It's just using a pot of the company's own money. This, again, is to make sure that you haven't made a mistake somewhere. We continued to see performance that was actually diversifying from what we had, from our other models, and making good returns.

Eventually, this whole process of test trading came to its conclusion. We had a very thorough peer review done. We were, at that point, able to put these things into the client portfolio. The only reason they're in the client portfolio is because they've earned a place there. It's not because we had a preconceived idea that we wanted to be the machine learning people. We started doing this back in 2012. It's taken a long time to get to where we are now. And the journey is still some way to go. I'm going to show you some of what we've done and some of the challenges we had to overcome. I'm going to give you a couple of specific examples. Just so that you know this is a material part of the business for us. It's not just some kind of curiosity. We are trading hundreds of millions of dollars of client money in our portfolios, using, explicitly, machine learning models.

These things have extended from where we first started, which was trying to capture new diversifying strategies. That was where we started. We then moved on to look at can we relearn the things that we've been doing for years using machine learning. One of the things I'm going to show you today is what you can discover that's new in momentum trading by using this powerful toolbox. The final thing we've been doing over - just over the last year, is look at how we can use machine learning, not within the strategy development, but how we go about actually using it within the contexts of trade execution. Some of the things that you've seen downstairs, when Richard presented to you at the big screen earlier on, those are controlled by machine learning systems that are live. I will show you some of those as well.

I've only got 15 minutes to tell you about over five years of work, so this is not going to be a complete treatment. If anyone wants to talk to me afterwards, I'm very happy to answer questions. I'm very happy to show you more. The first thing that makes us unique compared to a lot of competitors is that we have a research lab, based in Oxford, with our staff in it, that's embedded within Oxford University. No other fund management firm has that. We call our commercial research lab there the Man Research Laboratory. It's in the same physical building as the Oxford Man Institute. The Oxford Man Institute was opened 10 years ago and from - its focus has moved around within the space of technical, science and trading areas during that time. Its focus since August 2016 has been on machine learning.

So we have a very, very strong collaboration there with a group of internationally-renowned experts in the machine learning space. That's an incredibly useful resource for us, because if we want to fast track how we go from a concept and an idea, through to actually something that is implemented, we've got a group of experts on hand. We see them every day. We're in the same building. We have the same coffee machine. It's probably not as nice as the coffee downstairs here. It's the same canteen. Everyone is mixing. That is a unique resource and it's been incredibly useful for helping us in the journey of going from zero experience in machine learning to where we are now.

Now, just to give you some context about machine learning, you see this stuff appearing all over the place. The media is obsessed by it. Whether it's to do with self-driving cars or beating the world Go champion, there's loads and loads of activity out there using machine learning. Many of these things are quite revolutionary, but they're not typically in the finance space. If you look at, more specifically, to our intentions, then there is activity out there. This is the headline from last November. It's referring to a report that was put out by the FSB. Most of the comments here see this as some sort of destabilising risk. There are a few comments at the end that - it talks about some segments of the financial system are actively employing artificial intelligence and machine learning to find signals for higher and uncorrelated returns. We're doing that.

Investigating using AI tools to assess market impact. That's again something we're doing. That's controlling part of the execution world I was talking about. I don't know whether this is in reference to us, but, certainly, we are active in those areas. Now, what is machine learning, first of all? It's a kind of messy collection of algorithms that are about processing data to find structure in it, but it doesn't sit within statistics, or engineering, or computer science. It's a big amalgam of lots of different bits. The thing that distinguishes them from the more mainstream statistical tools is that these things have to be smart enough that they can identify structure in data without explicitly being told what to look for. They have to identify what sort of nature the model is and assess the complexity of that all by themselves. That's what distinguishes them from the more traditional statistical modelling.

It's very broad, many different approaches and toolsets. It's got a long history of development as well. This stuff, the media is obsessed with it currently, but it's been going on since at least the 1950s in one way or another. Now, the only thing I want to draw your attention to on this slide are these last three points here. Pretty much everything we do falls into one of those categories. It's about modelling complex relationships in data. What I mean by that, is something that's got more intrinsic structure than the usual type of linear models people come up with. It's about discovering new factors. One of the promises of deep learning - this is an area of machine learning - is you don't have to say what type of factors you're interested in. You don't have to define, well, I think momentum should work like this, go and find it in the data. It has to come up with its own predictors. That's been a very important area of growth for us.

The final thing we use it for is unlocking new datasets. We've already heard about the flow of datasets coming in. If we can automate that, make it more of an industrialised process, rather than a bespoke process with a human involved in it, then we can turn those datasets around much more quickly. Machine learning allows you to do that simple broad pass through things, to find some key pointers for where there may be some new predictors that we haven't come up with before. Now, why is it hard? Most of these techniques aren't developed in finance. A lot of them come from computer vision. A very famous example of computer vision problems is reading hand-written digits.

This is the first challenge that we've had to overcome, is getting these techniques, which are particularly good for these types of vision problems, to work in financial data. The main problem you first encounter is something to do with signal to noise ratio. Now, we can all look at that, and see that those are nice crisply-written white characters on a black background. There's different handwriting there. That's what the computer is doing. But there is no ambiguity about where the edge of the line is, or the edge of the signal. All of these things work by digitising, over some lattice, what is observed in the handwritten digit, and then assessing that against some sort of matching template. In this case, it's pretty obviously a three. Now, there's no ambiguity here in what the signal is, whereas, financial data just doesn't look like that.

These techniques are designed specifically for environments where you don't have this signal-to-noise ratio challenge. If you want to think what the numerical digits would look like in our world, it'd be something more like this, where the character is just barely distinguishable from the background noise.

In fact, it's even worse than this, because we're interested in prediction. You actually would have this character being just visible above the background noise, but you've only seen this part of it so far. You've got to work out whether it's going to be a two, a three, or an eight. If that's all you see against the background of noise, it's very different to that original handwritten character problem. So that's the first challenge. The second one is that finance data just isn't stationary. Threes are always written this way around, whereas, finance data, things can change. The correlation between stocks is positive sometimes, it's negative at other times. The timescales at which things pan out can also change. This non-stationarity again makes it very difficult to take these techniques from computer vision, or other areas, and apply them to finance data. That's what we've had to be doing.

There is a third challenge, which I'm not going to say that much about. If these things are calibrated just by historical data, what are they going to do if it's just a black box that you can't understand, that sometimes tells you to be long or short with the market? What's it going to do when its confronted with data that's outside the range of anything it's seen before? Unless you've got a very good knowledge of what's going on inside here, you shouldn't be trading other people's money on the basis of that model, if you don't know what it's going to do when it's confronted by data it's not seen before. And it will be at some point. Those are the challenges that we've had to overcome. Now, in terms of the spread of all the different things we're working on, I've mentioned non-linear signal combination. That's been very successful. That's gone into the client portfolio.

We've been using neural networks, or deep learning, really mimicking the kind of approaches that have been so successful in computer vision. That is successful. It's in the client portfolio. On the pure data structure discovery side, we've been using various techniques there, including text processing, natural language processing. That has gone into the client portfolio. The last thing is an area of re-enforcement learning, which is a way you learn about your environment through actually exploring it. That's what we've been doing in the trade execution world. I'll give some examples of those two things. Now, I said at the start, where we originally worked was on looking at things which were deliberately not correlated with momentum. I'm not going to cover that now. If you want to talk about that, we'll do it afterwards.

Just to finish off, please remember that machine learning doesn't always get it right. People assume it's got some sort of almost magic-like properties. That's not the case. It will, and does, get things wrong, even in the context of computer vision, where it's been tremendously successful. It correctly classifies these things as a centipede, an electric guitar and a starfish. That's great. But it also thinks that that's a centipede, an electric guitar and a starfish. These things can be easily confused. Here are some more examples of them getting confused. Just because it says machine learning, don't assume it's some magic system. It will always have a chance, and a reasonable chance, of getting it wrong. That's in computer vision, where it's an essentially solved problem, but not even in computer vision is the question that clear sometimes. These are my favourite ones - surprisingly difficult.

I'm going to finish off there. That's all I really want to say about machine learning. I'm happy to talk more. These things are incredibly useful, incredibly powerful. They're hard to implement, and do it well, because they really have their genesis in an area that's nothing to do with finance. The main hurdle of applying them in our world is overcoming those different design criteria, from where they come from to where we're hoping to apply them now. Thank you very much.

Luke Ellis

Anyone got a question?

Unidentified Speaker

Hi there. Just a quick question. Earlier today Giuliana put up a slide showing us the cumulative extra return from new models at AHL over the last few years. Do you happen to know if machine learning is in her slide? And if whether it is or not, can you give us some idea of how your machine learning has helped? Thank you.

Luke Ellis

It is in her slide. We don't separate out machine learning answers from other answers for clients or anybody. One of the things that I do think is important is, we have avoided, and we will continue to avoid, having a machine learning-only fund. One of the great things in machine learning is, if you've got a bunch of very smart people who understand what's going on, they don't get over excited when you get a period of amazingly good performance. And they don't try and shoot themselves when you get a period of bad performance. The nature of the models is that if you offer them solely in that form, to a client, it's really hard for the client, because they don't really understand the detail. When you get a drawdown, when you're saying the model - someone Sandy I think, someone earlier mentioned the fact that, as you'd have seen it, that one of the machine learning models in February kept trying to buy the dip.

And by the end of February it looked pretty stupid, because it kept trying to buy the dip, and it was wrong. That is the type of thing that is difficult for clients to be comfortable with. It's much better employed - these are techniques that are employed within other funds, so it's a core component within Dimension, but, actually, across a number of things. As Anthony mentioned, in execution everywhere, techniques are used in Numeric. They're even used now in GLG in a variety of ways.

Unidentified Speaker

How much scope do you have for human override in terms of when you're deploying machine learning?

Anthony Ledford

In terms of the way these things are deployed, they are deployed in exactly the same way as any of the other models within our portfolio. We will look at the performance that is generated on a standalone strategy basis. We have thresholds of performance which are set completely objectively, rather than being done on a discretionary basis, as to what level of performance starts to look unusual, according to the amount of volatility you're running and the strategy design. There is no difference as to how we run the machine learning models compared to the other models in our portfolio. We probably have more scrutiny on them, because they're newer, and we've had to develop bespoke tools for being able to understand what kind of trading the model is going to actually be doing. Other than that, they are identically the same as any of the other models and controlled by our standard risk process.

Luke Ellis

. Thank you, everybody. We've had a bit of a rush through to try to get a lot of different ideas in front of you all. Hopefully, you didn't mind the rush through. Hopefully, the passion that everybody has got for technology and what it can do in our business, for quant techniques, what they do in our business, has come across. Every one of the speakers was complaining about it's impossible to talk in 15 or 20 minutes, and couldn't they have an hour. Anthony, you can watch on video quite a lot if you look up AHL. It explains on YouTube. There's even some new ones, which is good for those who've looked at it before.

You've seen a number of AHL people today rather than a fully-diversified group from the different areas. That's proximity as much as anything else. Shanta came over, but we didn't want to fly too many people over from Boston. The techniques are being used across the firm and, hopefully, that came across, and in a wide variety of different ways. The more you can do to generate either better alpha, or to harvest more of the alpha through better execution, better portfolio construction, better risk management, both the more returns you can deliver to clients and, yes, the more money you can run for the same level of risk and return for the client. So it benefits all of us from that point of view. We finished reasonably on time. Thank you all for your time.

[End]