Verus®

How to make manager research decisions

April 2021

IAN TONER, CFA
Chief Investment Officer

MARIANNE FEELEY, CFA
Managing Director | Public Markets

Executive summary

How you ask a question matters. It determines the tools you use to answer the question and what a useful answer looks like. Investment professionals tend to ask questions about manager selection as though they were trying to ***prove*** *that the manager they were researching was going to add value. This makes reaching a meaningful answer almost impossible for a host of practical and theoretical reasons. Once we recognize that manager research can never be about proof, but is in fact an exercise of making* ***calculated balance-of-probability assessments with limited information****, and then* ***updating those assessments*** *as new information comes in, we can think more clearly. This recognition helps us in a number of ways. It clarifies how we should do research and how we should make manager hiring decisions. Most important, this change can help us improve our decision making: the grand tradition of investors buying too late and selling too late is based at least in part in the feeling that we need to "prove" the skill (or lack thereof) of a manager, and of course by the time the investor feels that "proof" has been achieved the story has moved on, and the "proof" is either no longer true, or is too late to implement. Changing our mindset can help us move more swiftly when there is meaningful new information while stopping us from getting distracted by information that seems important but which has little real value.*

Introduction: Gurus, experts & normal people

The investment industry is full of moderately numerate people: able to do most types of investment math on a day-to-day basis, but not technical

enough to dig deeply into advanced mathematical concepts. The numbers in finance are observed and plugged into tools – and generally people spend only enough time studying those tools to understand how to use them, rather than actually understanding the real underlying mathematical concepts and assumptions underlying the tool being used.

A useful analogy is automobiles. The relationship people have with automobiles falls into three categories.

— The first group is the small category of people who bleed motor oil. They really understand the mechanics and electronics of the modern car, can diagnose and fix any problem, and can probably build a fully functioning vehicle from scratch in their garage during their free time.

— The second group is much larger: they can talk about automobiles knowledgeably enough to sound convincing to most other people, and can probably follow most of the explanations of the first group when they take their car in for service, but if challenged to describe anything more than the most basic mechanics of a modern engine would likely rapidly change the subject. This second group are at least knowledgeable enough to be able to do some basic vehicle maintenance if they have to and to have some insight into the right type of vehicle that someone might need if they were considering a purchase.

— The third group simply regards their car as a way to get around, focusing on comfort, gas mileage and cost of ownership. They know, in a general sense, something about how the vehicle works, but limit the amount of time and energy they spend on learning about all things automotive to no more than the amount needed to be able to make an adequate decision about which to buy when that time comes around.

We can call the first group the gurus, the second group the experts, and the third group normal. There is nothing wrong with whichever group you find yourself in – but it does matter that you know what you know, and where you are bluffing when you are talking about technical details.

The same thing applies to the use of tools in finance. There are three similar groups in the intersection of math and finance.

— The first group, the gurus, really understands both the mathematical ideas being used to explain markets and investments and their limitations. In particular they understand the basic embedded assumptions that are the basis of the theory that justifies the models they use, and they know which of these assumptions matter for the calculations and which can be ignored without affecting the result, even if they don't strictly fit the situation. These people tend to be behind the scenes, working on analysis, rather than spending much time with clients – often their communication skills are not as developed as their mathematical understanding.

— The second group, the experts, are the majority of investment professionals. Their expertise is typically broad, not deep: they have to know about a vast array of topics in some detail and this means that they will likely only have deep expertise in some areas. Some of their most important skills are around client management, for example, as the heart of their job in many cases is focused on helping their clients understand their goals and the risks they are prepared to run to meet those goals. They will generally not question the tools they are given in any detail, and will simply assume that the approach supported by those tools makes sense.

— The third group we will again call normal. These are similarly people who are primarily consumers of financial advice, although of course some non-professionals spend enough time and focus to become experts or indeed gurus. The normal group, however, is forced to make decisions about investments whether on their own behalf or in a role as part of an investment committee, and do so based on the advice and using the output of tools given to them by their advisors, who are mostly in the second group, the experts[1]. This third group are often wise, bringing both intelligence and an external perspective to the conversation, but they will still generally operate on the assumptions given them by the advisors they work with.

The distinction between these three groups matters for one simple reason: while most of the time the tools that the experts employ are being used for the right purposes and in ways that are helpful, sometimes the assumptions underlying those tools simply do not work, and more importantly are not helpful for making sensible decisions. How we approach manager research decisions is one of the places where we believe there is this disconnection. In fact, the desire by the expert group to use what seem on the surface to be the most robust mathematical methods to make these decisions turns out to be an excellent example of that disconnection. More "rigor" is only better if the rigor is coming from the right discipline – and in this case it may not be.

Two ways of dealing with uncertainty: "proving it" or "improving it"

There are two quite distinct ways that humans can go about dealing with uncertainty about things they see around them. Both are entirely reasonable, but they fit different types of situations, and apply to different ways of thinking about decision making.

— The first approach assumes that there is an underlying "true and constant" signal that exists but that we cannot see directly, and that our job is to observe data with a goal of identifying the point where we can "prove" that the signal exists. When we observe samples of information we are asking the question "what is the probability that the thing we are observing is due to something other than random chance" – our goal is to be close to certain that the "signal" is causing an effect that is not just random. We use terms like p-values and confidence intervals in this type of work. We can call this type of thinking "Proving It" thinking.

— The second approach focuses more on decision making. It assumes that we have a set of information, experience, beliefs and knowledge about a particular situation based on which we make certain assumptions about the world, and then asks questions about how to update those assumptions given new information that we observe. The focus is less on "proving" things, and more on the process of updating a set of beliefs we have already created – called our "prior". When we observe new information we are asking the question "how can I update my prior given this new information" – our goal is to help improve our ability to make decisions about the future under uncertain conditions, given new information. We can call this type of thinking "Improving It" thinking.

What are the key differences between these approaches? Probably the most important is that the "Proving It" approach depends on the idea that there is an underlying hypothesis that is constant and that can be proved true by collecting random samples of data of observations. It fits well when doing science experiments about natural physical laws, or testing the effects on crop yields of different types of fertilizers, where there is a lot of data available, and where the underlying effect they are trying to prove is unchanging. It does not rely on the beliefs of the person doing the sampling, although those beliefs might help them develop the hypothesis they are trying to test.

The "Improving It" approach, on the other hand, does not focus on testing a constant underlying hypothesis with random samples of data, but focuses instead on understanding the degree to which new information should cause the experimenter to update their beliefs about the world. Although there is mathematical rigor underlying it, the basic approach involved can be used to guide thinking even in very qualitative environments. It is an iterative process: creating a "prior" view of the world, making observations, then updating that "prior" appropriately based on the observations. This is well suited to situations where there is very limited data.

It is fairly clear that these are two entirely appropriate ways of thinking about and dealing with uncertainty – but it is also fairly clear that they are very different approaches. It may also be unsurprising that there is a long history of people who approach problems using one mindset having violent disagreements with people who approach problems using the other mindset.[2] It is also important to note that most of those disagreements really boil down to arguments over which tool is appropriate for which job, not whether one approach is right and the other wrong. What is true, though, is that many of the statistical approaches and tools that business schools teach are focused around one of these approaches – the "Proving It" approach. It is also true that the output of "Proving It" tools can look more mathematically impressive than the alternative, even though the underlying rigor of the two approaches is the same, and mathematically impressive output is a valuable sales tool in the financial advice profession, even where better approaches may be available.

That, then, is the next thing we will turn to: which approach fits the kinds of questions we are asking when we do manager research and manager selection?

Is manager research a "proving it" or an "improving it" process?

When we talk about investment managers, the language that we use, and the analysis that we perform, sounds very much like we are talking about a "Proving It" process. There are a number of pieces of information that investors look at to try and perform analysis. Investors generally expect a certain number of years of a performance track record from the manager, and then perform detailed analysis of the return stream generated by that track record. Although this analysis usually stops short of formal statistical metrics, the underlying ideas behind the conversation usually sound a lot like a "Proving It" process. Indeed, most of the data analysis that we perform (analysis of tracking error, beta, drawdown statistics, upside/downside capture ratios and so on) is at least implicitly based on the idea that the return stream is statistically meaningful[3].

This seems like a very bad fit to the reality of investment manager research. Even in the most systematic investment process there is significant change over time in a way that should make us doubt the stability of the data. Imagine that the minimum period of data that we would need to reach any form of certainty in this type of analysis is maybe five to seven years, with ten or more being preferred[4]. This applies across a number of dimensions:

— First, think about the investment management company. Has the management structure stayed the same? What about the compensation and bonus structure? How has ownership changed over the period covered, and how has that affected the alignment of the team who are managing the money? Are there other lines of business or other products that the company runs – if so, have those provided financial support for the team running the product, or have resources had to be drained from the product team to fill gaps opened elsewhere? Is the company a target for acquisition (or even being talked of as such), or is it acquiring other firms itself?

— Next, what about the investment team. How has that team changed over the five, seven, ten or more years? Are the people the same people, in the same roles? Even if there have been no changes in titles or roles, how have the things those people have done changed over time? What succession plans have been put in place and how has that affected how people work and behave?

— Next, about the people themselves. Nobody stays the same over ten or fifteen years – certainly the main things about an individual stay similar, but over that period they are likely to change, learn and develop. Some of those changes will be good, other less so, but every day each person in the team will learn some things, forget some others, and change some of their views, even if only a little. How have those changes affected the way they do their job? How is the manager approaching managing money at 55, and how does it differ from the way she did so at 45 or 35?

— Next, think about the technology and tools available to the manager in running money, and also available to their competitors. How have these changed the day to day of their process, and what parts of their edge have been helped or hindered by those changes?

— Next, think about competitors more generally. How have they changed how they invest – how many others have grasped some of the insights that the manager originally had that justified the product in the first place? What has the manager learned from their competitors?

— Market structure and nature may have changed too – how have those changes impacted the manager and their ability to add value? What about the nature of the global economy, companies, industries and capital markets: how have these changes affected the way the manager can add value?

This list could go on for some time, and each issue suggests two more. The point being made here is not that the data we have about investment products is useless: it is nothing of the sort. However, what it is absolutely not is a robust, continuous data set. Instead, it is a collection of time-series data representing a changing process, by people that change, in a changing market with changing technology in changing economic conditions[5]. Using this data to "prove" the manager has skill seems unlikely: surely using tools and thinking based around the idea that we are trying to do that type of "proof" seems unhelpful.

When we turn to the "Improving It" approach things seem to fit much better. In reality, good manager research analysts do not simply look at certain facts and then perform analysis of those facts: instead they look for certain characteristics of products and then draw conclusions from the presence of those characteristics. They draw a conclusion as to whether each particular product is likely to be effective based on those things that they look for. They then, on an ongoing basis, perform a series of observations, updating their conclusion based on the relevance of the information observed and the signal given by that information. Each decision they take is based on a mixture of qualitative and quantitative measures, and they are balanced based on their likely explanatory power. At no point is there "proof" of the skill of the manager: we are instead in the world of drawing conditional conclusions, and accepting the risks of failure that come from that process.

This "Improving It" approach seems to fit much better both with the task we are trying to perform, and with the actual data and information that we have available. In the next section we will look more at what this means for the manager research process, why this insight calls us to look for specific things in that process, and how we can use this approach to help investors make more effective manager hire and fire decisions.

Manager research: The "improving it" approach

The core elements of the "Improving It" approach comes down to three steps:

— Creating a prior: using all of the information available to us, in the context of our experience and past history, to come to a conclusion about the world. In this case our prior will be one of the following:

- I have good confidence that this product will do better than most other products over a full cycle
- I am not sure how this product will do over a full cycle
- I have good confidence that this product will do worse than most other products over a full cycle

— Making observations of new facts and information that we think might be relevant to the priors we have created

— Updating the prior based on new facts and information

The core of all of this is a simple question: what facts and questions help us create and update this prior? A sensible manager research process should focus on those, rather than spending time on other topics. What would this look like?

BUILDING A PRIOR

It seems clear that we should focus our research on the things which we believe are going to be material in helping us establish a belief in the likely success of a product, whether or not those things can be measured exactly, and whether or not researching them is convenient. Simply concentrating on easily available data that can be used to produce what appear to be sophisticated charts is likely to be ineffective. While the charts may appear impressive, the only thing that matters is ***whether they are relevant and helpful in creating or updating our belief in the product.*** This implies that we should be focusing on looking for particular characteristics of products or managers, spending time on them and less time elsewhere.

In the Topic of Interest paper published in August 2020 "AEIOU>PPPPP" we concentrated on discussing the key principles for manager research: the things that we look for in manager products. They are as follows:

— Alignment

— Edge

— Implementation

— Optimal use of risk

— Understandable performance

We contrasted this with the more typical approach – defined as "the Ps" – which looks at certain things like Philosophy, People, Process, Performance and Parent. We differentiated these by clarifying that the Vowels represented things we were looking for, while the Ps were simply things that were looked at.

Now that we are becoming more formal with this "Improving It" framework the reasoning behind this becomes much clearer. We focus on the Vowels during the manager research process because those are the most relevant things to look for when constructing a prior: they may not all be highly quantitative, but long experience of the research process tells us that managers with exceptional showing in these categories are well placed to produce good outcomes. In other words, these are the key things to focus on when creating a prior belief that the product is worth backing.

For example, we can consider the Alignment topic from our "Vowels". This is one of the single most important things that can make or break an investment organization: create a culture and structure where the key staff that are driving outcomes are appropriately aligned, incentivized and compensated, and that the more junior colleagues have line of sight to equivalent treatment in their future. This builds long-term stability into the organization, and therefore the outcomes the organization generates. Fail to do so, and you significantly increase the probability of disruption of the organization, and therefore likely the outcomes in the future. Trying to "measure this scientifically" is almost impossible, but gaining comfort on this topic is vital, as material positive or negative changes in this area are likely to have a big impact on outcomes.

What is notable is that none of the Vowels are formally defined in a numerical way. We are not looking directly ***at specific metrics:*** instead we are looking ***for specific characteristics***, using both qualitative and quantitative measures to draw a reasoned conclusion as to the strength of showing in that category by the manager. The conclusion is a rating – good or bad – and that rating doesn't represent an entirely scientific conclusion, but instead ***a well-considered opinion*** that a product is ***likely to do well or poorly***. We then follow the appropriate approval process to rate the product and move on. The exact rating system we use to describe the outcome can vary – there is no one perfect way to describe how you rate managers – but however you describe the conclusion, what matters is how that conclusion is reached and updated.

OBSERVING NEW DATA

This same approach carries over to the next step of the appropriate process. Rather than simply focus on the performance outcomes, we want to look more broadly. Having done the work to identify that the five Vowels are the core drivers that are likely to be indicative of managers that can produce good outcomes, we want to focus on things that affect them. Simple performance updates are, if not useless, at least only rarely informative. Instead we focus on changes in alignment, challenges to the investment edge claimed, changes or failures in the implementation approach, use or measurement of risk that seems sub-optimal relative to the edge claimed, and performance that seems difficult to explain given the other principles.

The key to this collection and observation of new data is focus. The temptation is always to perform detailed analysis of data that looks sophisticated but which has little to no impact on

the principles outlined in the Vowels. Instead, we want to collect and consider new data in places and in ways that are likely to be relevant to the things that we believe are actually going to drive investment outcomes.

UPDATING OUR PRIOR

The final step of our process then involves us updating our prior. If the world were a less inconveniently nuanced place this updating approach could be robustly mathematical, but the reality of manager research suggests that this would be both too complicated and likely overkill[6]. Instead, we can simply set up a clear table of the principles and how the new information helps us update our conclusions on each of them, then assess how this causes the net conclusion to change or to stay the same.

Advantages of "improving it" in manager research

What are the advantages we get from being clear about the decision process we are adopting? Those advantages fall into four categories:

- **Clarity:** Because we focus on a simple iterative process, and because we are clear that we are adjusting a prior, and that the prior is based on qualitative and quantitative assessments of each of the principles, we have the advantage of clarity. We know what we are looking for each time we reconsider and adjust our prior, and we know what to ignore.

- **Appropriateness:** Because we are clear that the decision we are making is not one involving proof, but one involving judgement, we can approach it in an appropriate way, only spending time with data that is relevant, and only performing analysis that is appropriate for that type of decision making process.

- **Timeliness:** Because we can focus on regularly updating our prior, rather than attempting to prove that the manager has skill or that they do not, we can react more quickly and effectively if needed. If all of the decisions we make involve judgement rather than proof, identifying a small piece of information that fundamentally shifts that judgement can prompt rapid action.

- **Diversity:** Because the focus of this approach is on creating sensible priors using a range of inputs, we are not tied to particular data sets, including length of track record. Instead we can think about the prior creation and updating process more broadly, and can where appropriate reach positive conclusions about newer and smaller firms where reasonable basis exists. This helps nudge the conversation in the direction of increasing openness to diverse and emerging managers.

So, do I buy, sell or hold?

Although we have touched on this final key question throughout this paper, it is worth finishing with a direct response to the question: how do I decide whether to buy, or more importantly sell, a manager? The simple answer is this: to recognize that the drivers of this question are focused on the "Vowels", which are the things which influence the likely long term outcomes from the product. When the facts observed around those factors cause an update to the prior beliefs about the product's ability to add value, you should divest – but not otherwise. Many of the things that can be observed – particularly performance – are only lightly relevant to that updating process: many of the things that are relevant are hard to observe, or at least hard to quantify. What this means is that these decisions will ***always*** be balance of probability ones, and the best of them will happen before there is actual ***proof*** that the product no longer deserves the investor's trust. Understanding and acknowledging the discomfort that this type of decision will always cause is important: framing an investment decision making process around that type of decision is likely to produce better decisions. Some will be correct, while others will be wrong, but on balance this approach is likely to be more effective, and to pull decisions into a more helpful structure. Waiting for proof may be like waiting for Godot – the proof you are looking for may never come, and some may find the process of waiting quite uncomfortable. We have, in the appendix to this paper, some worked examples of how this process can actually work, along with a simple tool that can be used as a ready-reckoner to help in that process.

Conclusion

Throughout this paper our emphasis has been on changing the way that investors think about making decisions around investment management products. We believe that by focusing on the "Vowels" we clarify our decision inputs, and by structuring the decision appropriately around creating a belief about the product, and then updating that appropriately as new information becomes available, investors can make better, and clearer, decisions. We also present in the Appendix a simple schematic investors can use as a thought aid to help guide those conversations and decisions and give some examples of how this approach might be followed in fact-situations approximated on real life. Actively concentrating on what type of decisions are being made can help investors focus on the things that may help them make those decisions more effectively. It can also make the fundamental job of the trustee – making difficult decisions under uncertainty – more comfortable.

APPENDIX: WORKED EXAMPLES

To clarify the way in which this thought process can actually be used, we include in this Appendix both a format that can be effective as a ready-reckoner and some worked examples of typical situations that might occur in the real world.

Worked examples: A decision format

We can now introduce a format which describes the thought process that best fits manager assessment work. This format is expressly designed to be extremely simple: a short table, updated with new information and then reassessed. The table is designed so that it could be used as a summary of all of the work performed, and to describe the framework within which we draw our conclusions – it can be thought of as the "front page"[7] of a longer body of work, likely based on many hours of research and analysis. This simple table structure helps clarify that, at heart, the final decision we are making is always an "Improving It" decision around beliefs rather than a "Proving It" decision that comes mechanically from analysis.

This simple table is useful for purposes of this paper, and can be used by researchers or Boards trying to think about the products that they are assessing, or more importantly, considering hiring or firing. Anyone coming to a conclusion about a product should be able to fill in this table in less than two minutes as a "gut check" on the conclusion they have reached – if they have the level of knowledge and understanding about the product that they need to come to a reasoned decision. If they find filling in the table difficult it can indicate both that they may have more questions to ask, and which questions should be asked. This table is not designed as a formal part of a standard product description report, or as a formal part of a search report, but instead as a ready-reckoner to help analysts or decision makers guide their thinking process. A really important feature of the table is that the individual inputs are not designed to "add up" in a mechanistic way. Not all inputs have the same weight, and the appropriate weight for each input may vary by product based on the firm and product involved, so there is no single way to cumulate the inputs. Instead, the conclusion drawn is the result of all of the issues taken in context of the product being assessed.

We begin with how we summarize our prior.

A	E	I	O	U
+	+	+		
Conclusion: We believe that this product is likely to provide good results over the long term				

Each of our principles is represented by the appropriate vowel – Alignment, Edge, Implementation, Optimal use of Risk, Understandable Performance. Under each vowel is a symbol representing the information and beliefs on the basis of which we reach our prior. All of the extensive work done to back up the conclusion is in the files on which this conclusion summary is based: this single box simply represents the summation of all of that complex

work. In this case, then, the manager has described and demonstrated alignment characteristics that we believe are good, have described and demonstrated their investment edge in a way that we believe is good, and has described and demonstrated implementation characteristics in a way that we believe is good in the context of the edge being claimed. The manager has not described or demonstrated characteristics around the optimal use of risk or understandable performance in a way which has caused us to create a belief that they have enough substantive strength in those places to influence our overall opinion of the product either positively or negatively.

We then update our prior using observed information. This observation will likely be in some detail, and when written up may take a number of pages, or some complex analysis – it can be fully quantitative, fully qualitative or a mixture. We carefully consider each piece of information, think about which of the principles it applies to, determine the degree to which this information taken together should potentially cause us to adjust our prior, and then update our conclusion accordingly.

Principle	A	E	I	O	U
Inputs to Prior	+	+	+		
Observations	All junior members of PM team granted equity		Significant investment in technology on trading desks to improve speed of execution of trades capturing very short-term mispricing	Recent position size band failings, not addressed well in conversation with risk team	Good performance relative and absolute during last six months. This is surprising given manager claimed edge and approach
Changed Inputs for Updated Prior	+	+	+	-	-
Conclusion – Updated Prior	We continue to believe that this product is likely to provide good results over the long term. We will do further research into whether the source of this good performance should cause concern about whether the manager is losing focus on the approach to investment that generates their claimed edge.				

This simple table captures the entire process in one simple figure. We begin with the inputs that helped us create the previous prior. We then perform observations, which are summarized in text form in the next row. These observations are turned into inputs for the updating of the prior. Finally we reach our conclusion – our updated prior, which we can then use and which becomes the starting point of the next iteration of this process. The table we use here can help guide the process, although it is of course not the process itself, and each board or investment team can use a format that works for them. What matters is the content and the process, and in particular the clear understanding of the "Improving It" nature of the decision being made.

In the example above there was no change in our view. We observed some data points that were likely to reinforce or strengthen our previous view, and some that were likely to raise concerns. We identified some performance data that sat uncomfortably with the claims the manager made for how they generate performance – that caused concerns even though the performance itself was strong. We determined that the new facts that we observed were not different enough in important ways to justify updating the priors we had created, putting more weight on the inputs that were relevant to the particular circumstance and less on those that were less so, and so we reaffirmed that prior. We then pointed the way ahead to the next iteration, identifying what further work we would undertake in an area where we identified that adverse information would be highly relevant to a reconsideration of the prior.

We can perform a similar simple thought process in a couple of other examples. We will cover two here. The first describes the process when initiating coverage of a new manager with limited track record. The second describes the decision to sell a manager. Both are drawn in part from real life fact situations, although we have changed enough details to make the manager non-identifiable, and the decisions described represent potential reactions to those changed fact situations.

Limited track record – Moving to hire

This first situation describes a manager that has been set up in the recent past. The portfolio manager has had a good history at other firms, although the move to a stand-alone structure means the manager is taking on broad P&L and management responsibility, not just the investment leadership role. The structure appears sound, and there are external shareholders providing financial support of the organization. The strategy being followed is one that has succeeded in the past, and the CIO has built what appears to be an effective small team to provide the needed capabilities to deliver the outcomes expected. The investment edge is clear, and the performance in the short time since the firm was established appears in line with expectations.

Principle	A	E	I	O	U
Observations	Team are equity owners, founder with majority of his personal wealth invested in product, firm manages only one product	Focus & consistency, with a culture of heterogeneity across the team. CIO is a strong and effective leader	New firm, but philosophy & approach has 15+ years history in other firms, key team members worked together elsewhere, good relationship with company management despite small size	Moderate concentration with highest concentration in highest conviction names	No material observations due to relatively short track record.
Inputs to Prior	+	++	+	+	
Conclusion	We believe that this product is likely to provide strong performance relative to peers over the full cycle. Ongoing observation will focus on the relationship between the tracking error derived from concentration and the expected alpha generated from the strategy, in particular with reference to the relationship between the expected drivers of return embedded in the long term growth approach at the heart of the claimed edge. We will also watch the management capabilities of the CIO/CEO.				

The research process to reach this conclusion took around a year from first contact, and focused by necessity on the key principles that gave us enough confidence in the team and product to reach a coherent positive prior despite small asset size and short performance track record.

Risk management concerns – Move to negative opinion

The second situation describes a firm where new work uncovers concerns about the organization which cause a re-evaluation of the balance of probabilities, even though those concerns have not yet caused significant pains in terms of investment outcomes. The organization concerned has a good reputation, with high quality investment professionals who are extremely knowledgeable and skilled in their specialist field. Led by an exceptionally well thought of portfolio manager, who also leads the firm, they had produced good results, although there had been occasional periods of road bumps. Discussions with the team, however, combined with an on-site visit, raised questions about the risk management systems and tools being used, and about whether the output from those risk tools was really being taken seriously, or acted more as theatre, providing justification for the investment team's views when they agreed, and being gently dismissed when they did not. In the space concerned this could have significant downside implications. These concerns caused the analysts to change their beliefs about the product, as they affected both Alignment and Optimal Use of Risk, and also called into question the topic of Understandable Performance.

Principle	A	E	I	O	U
Initial Observations Justifying Prior	Independent firm structure	Strong, experienced leader	Long track record of successful implementation and skilled team	High beta and tracking error due to strategy adopted, deemed appropriate on balance after discussion with manager	Long track record, with good success over the long period
Inputs to Prior	+	+	+		
Prior	Due to the long track record and the experience and reputation of the leadership and team of the firm the analyst reached the belief that the product was likely to provide good results relative to peers over the full cycle.				
Observations	Some concerns about relative power structures within the organization causing an imbalance between senior investors and risk management	No new information	No new information	Concerns over approach to risk management, risk tools being used, and the actual impact of risk inputs on position taking	Concerns that approach to risk creates materially lower reliability of long-term performance information, and less confidence in future stable outcomes
Changed Inputs for Updated Prior				-	?
Conclusion – Updated Prior	We no longer believe that this product is likely to provide good results relative to peers over the full cycle, although we recognize that many investors will continue to place confidence in the investment team, who have produced good results in the past. Investors could reasonably search for alternative products in this space.				

The concerns raised during this process were important enough to be a reasonable basis for a changed opinion. At the same time the new information was difficult to "prove", and it was also entirely reasonable for an investor to determine that nothing had changed in the product or firm, and that the concerns required more backing to cause a change in opinion. In this case further work validated the concerns, and subsequent performance issues appeared to be related to the issues raised here – the "Improving It" approach demonstrated its strength as a way to understand and triage qualitative information about investment and risk processes.

Notes & Disclosures

1 *Although many firms may have one or more gurus, the degree of focus on a small number of very detailed topics required of a guru means that they are often not well suited to broader client-facing roles, where breadth of insight has an important role, and can be more helpful for most problems, once a certain level of expertise is reached.*

2 *A simple search of the phrase Frequentist vs Bayesian (the "Proving It" approach is a major simplification of the thinking underlying the Frequentist mindset, while the "Improving It" approach is a major simplification of a Bayesian mindset) will prove that the ability of intelligent people to disagree with each other rudely about complicated topics is a constant*

3 *In the same way that you do not need to yell "I'm wearing a parachute" when jumping out of an airplane at altitude, as that assumption is baked into the activity itself, you do not usually need to state the implicit assumption that you are making that the data is a meaningful data set with a relatively stable signal in it, the existence of which you're trying to prove, and the nature of which you are trying to analyze.*

4 *In fact, the time needed to get true statistical proof is likely much more than this, but that simply stresses the point.*

5 *The famous line that one cannot step into the same river twice – because both you and the river have changed in the meantime – is apt.*

6 *We continue to dig gently on this topic however, as some increased formality of calculation could potentially give greater insights and might help improve outcomes.*

7 *Sometimes literal, sometimes metaphorical*

Verus®

800 Fifth Avenue, Suite 3900
Seattle, Washington 98104
206-622-3700
verusinvestments.com