

Oxford
Risk

Behavioural finance.
Applied.

Under the Microscope: 'Noise' and investment advice

February 2021

momentum
investments

Contents

About the author	1
Forward notes	2
Executive summary	6
Purpose	6
Findings	6
Forms of influence	6
Adviser archetypes	7
Decision prosthetics	7
Introduction	9
What is noise, and why should we care about it?	9
Survey design	12
Developing hypothetical client profiles	12
Limitations	13
How noisy are investment advisers' recommendations?	15
Suitable risk	15
Risk capacity	15
Financial personalities	16
Portfolio recommendations	18
Summary: measuring dispersion	19
What influences advice?	23
Predictive factors: signals along the advice journey	24
Inexplicable variation: noisy advice	29
Adviser archetypes	33
Cautious	34
Unsure	34
Risk-tolerance focused	36
Relaxed	36
Conclusion and recommendations	38
Noise vs bias	38
Possible explanations for inconsistencies	38
How to cancel noise	38

About the author

Greg is a specialist in applied behavioural finance, decision science, impact investing, and financial wellbeing.

He founded the banking world's first behavioural finance team at Barclays in 2006, which he led for a decade. In 2017 he joined Oxford Risk to lead the development of behavioural decision support software to help people make the best possible financial decisions.

Greg holds a PhD in Behavioural Decision Theory from Cambridge; has held academic affiliations at UCL, Imperial College, and Oxford; and is author of Behavioral Investment Management. Greg is also Chair of Sound and Music, the UK's national charity for new music, and the creator of Open Outcry, a 'reality opera' premiered in London in 2012, creating live performance from a functioning trading floor.



Greg Davies, PhD

Oxford Risk

Head: Behavioural Science

Forward notes

There is little doubt of the value that financial advisers add in getting clients to accumulate, protect and ultimately distribute their wealth. However, financial planning is a professional service just like medicine and law. These fields too clearly add tremendous value but are also prone to the inconsistencies inherent to any profession where human judgement sits at the centre.

Recognising the importance of client and adviser behaviour in reaching investment outcomes, Momentum Investments has initiated a landmark study in partnership with the prestigious Oxford Risk in the UK and South Africa's professional body, The Financial Planning Institute, to shed some light on these inconsistencies in the investment advice process which are termed 'noise'.

Managing 'noise' is not about giving every client the same answer, but rather in ensuring that drift en route to an investment goal is rooted in client circumstance and not the practitioner in question.

Paul Nixon, MBA CFP®

Momentum Investments
Head: Behavioural Finance



In the information age the value of advice has shifted from a transactional relationship of choosing the best product to a long-term relationship where the financial planner acts as a coach and mentor. For the financial planner to remain relevant in the future, they need to upgrade their human and technological skill sets. To devote time to the important role of behavioural coaching the financial planner will need to employ technology to perform routine tasks, use their technical skills to blend the results from the technology to a plan, and their human skills to help clients make the right decisions. This will be an important element in reducing 'noisy' advice.

The FPI, as part of the global financial planning profession through the Financial Planning Standards Board, is in the process of updating the Financial Planner competency profile to make sure that the profession remains relevant in the future. Studies such as this will only enhance our understanding of client behaviours, and how best we as financial planners can serve our clients.

David Kop, CFP®

Financial Planning Institute of Southern Africa
Head: Policy and Engagement



The research that Oxford Risk is doing in conjunction with Momentum Metropolitan is completely consistent with the outcome-based investing (OBI) philosophy that underpins the design and management of all Momentum Investments' funds and portfolios. 'Noise' is negative both from an advice and an investment management perspective - particularly when, as a company, we are focused on delivering the best possible solution for clients' needs.

Evan Gilbert, PhD


Momentum Investments
Senior Research Analyst




All clients would benefit from a better understanding of the balance between their required investment returns and the risk that they would need to take to achieve their objectives. It is very easy for clients to forget what they want to achieve when market movements and volatility only places the focus on risk. As an industry we should do everything in our power to assist clients and advisers to manage emotional responses to their investments better to ensure better outcomes.

Anton Swanepoel, CFP®

Momentum Financial Planning
Acting Head: Financial Planning





Executive
summary

Humans add value but are error prone and inconsistent. Where their inconsistencies are systematic, we should provide tools to combat them.

Executive summary

Purpose

The purpose of this research was to detect variation in investment advice, and to understand where it comes from. It investigated questions such as:

- For the same client, do advisers agree on the risk level they would recommend?
- Which factors – relevant or irrelevant – influence their advice? For example, do advisers project their own attitudes to risk onto their assessments of clients?
- How much of the variation in advice is seemingly random or due to inconsistency – in other words, how much is noise?

Findings

Given a set of relevant and irrelevant client information, including risk tolerance, advisers were asked to assess each client's risk capacity, composure (a behavioural trait), and knowledge and experience, and then use these to recommend a suitable risk level plus a high-level asset allocation.

Client circumstances were 'paired' to better understand these influences, eg two clients who were identical save for their risk tolerance.

- **Suitable risk level** – When presented with the same hypothetical client information, advisers gave remarkably different judgments on how much investment risk would be suitable.
- **Asset allocation** – High-level asset allocations were scattershot. Even where there was agreement on suitable risk levels, there was disagreement on what kind of portfolio would represent this.
- **Risk capacity** – Advisers were very divided when it came to considering a client's risk capacity. Only one of the pairs of similar clients showed any consensus.

- **Composure** – The clearest indicator of composure in the information given – the client's stated risk perception – displayed a weak correlation to adviser assessments of composure. There was a stronger link between appraisals of composure and risk capacity, which theoretically should be unrelated.
- **Knowledge and experience** – There was an aversion to extremes in assessments of both clients' investing knowledge and the degree to which clients would like to have control over their investments. However, there was at least consensus as to whether each client was lower or higher than average.

There was a large amount of variation in the recommended risk level that doesn't seem to be driven by the information advisers were given on each client case study. Entropy analysis revealed that overall adviser assessments were closer to totally random than totally consistent.

The largest influence was adviser judgements of risk capacity. Risk capacity should probably have the greatest effect on the recommended risk level. More should therefore be done to ensure consistency in its measurement, and protect it from undue influences.

Forms of influence

In understanding what affects advice, we can break down influence into four groups:

1. **Explicable and relevant** – Factors we can explain, and that should affect advice. Within this category, influences can be in the right or wrong direction (for example, risk tolerance correlating positively or negatively with recommended risk level); or should affect advice, but don't (eg risk tolerance changing, but recommended risk level not).
2. **Explicable, but irrelevant** – Factors we can explain, but that shouldn't affect advice. For example, the sex of the client, or the age of the adviser.

3. **Inexplicable, across advisers** – Cross-adviser idiosyncrasies that cause advice to differ in persistent ways, without full explanation.
4. **Inexplicable, within adviser** – Random errors or inconsistencies that cause advice to differ despite being given by the same adviser to the same set of client circumstances, due to eg their mood.

By 'noise', we most properly mean variation that cannot be reliably explained or predicted, ie category 4. Inexplicable effects that persist across advisers may be due to factors not covered in this survey.

Adviser archetypes

Undue influence exerted by adviser circumstances rather than the clients' can be divided into broadly four categories, based on personality clusters that determine broad adviser archetypes:

1. **Cautious** – Project own lower risk tolerance and composure onto clients; more likely to work with mass affluent than high-net-worth individuals (HNWIs).
2. **Unsure** – Less varied judgements, lower confidence in recommendations, more likely to be paid by commission only.
3. **Risk-tolerance focused** – More responsive to a client's risk tolerance.
4. **Relaxed** – Project own higher risk tolerance and composure onto clients; more likely to work with ultra high-net-worth individuals (UHNWIs) and have fewer clients.

The 'Unsure' group are particularly interesting in that they are akin to archers firing off arrows in roughly the same direction every time, despite dealing with a moving target. Hesitant as to whether they'll hit the bullseye, they avoid recommending anything extreme, even where circumstances would suggest it were suitable. While influences based on the adviser's situation are unhelpful, they explain only a minor amount of the variation.

Decision prosthetics

Much like a roleplay, this exercise was obviously artificial, and not a flawless representation of the way adviser-client relationships, nor advice processes, exist in real life. This artificiality limits the weight to attach to any conclusions.

However, just as roleplays, however artificial, improve real-life performance, so should this noise audit. The real world in which advisers operate is messy, and full of potentially unhelpful influences. Mere knowledge of this messiness is not sufficient protection.

And while this may be a relatively simple model, when almost one in five advisers recommend more risk to a client who has lower risk tolerance than another who is identical in every other way, it's clear that more should be done to cancel out the noise.

This is not about removing the adviser's art, reducing the entire advisory process into an algorithm. It's about identifying high-level effects and minimising unhelpful influences, to greater allow the adviser's true worth to be heard above the noise.

Humans add value but are error prone and inconsistent. Where their inconsistencies are systematic, we should provide tools to combat them. Doctors still use checklists. Top sportsmen still use coaches. Chefs still use scales. And in financial advice, decision prosthetics can make advisers more consistently the best versions of themselves.

By 'noise', we most properly mean variation that cannot be reliably explained or predicted ...

Introduction

Introduction

What is noise, and why should we care about it?

The most suitable investment solution should depend on the client, not the adviser recommending it.

Consistency of advice is a crucial concern for any firm. If what is deemed suitable for a client can differ depending not only on which adviser within a firm they speak to, but also on the prevailing mood of a particular adviser, then that firm has a problem. Especially when we remember that advice isn't a single event, but an ongoing relationship, and that the regulations care not whether you get it right on average, but whether you get it right for each individual.

The two main sources of inconsistency in advisory processes are an overreliance on humans, and the heavily front-loaded nature of suitability assessments. In addition, there is frequently some – often reasonable – disagreement about how much a particular feature of a client's situation should influence the suitable solution. For example, some advisers may give the client's long-term risk tolerance a higher weight in determining the best investment solution, whereas others may place more emphasis on their current financial circumstances.

Humans are wonderful at many things. But they are inefficient and unreliable decision makers, especially where many moving parts are involved – as in risk capacity. Humans are prone to 'noisy' errors – unduly influenced by irrelevant factors, such as their current mood, the time since their last meal, and the weather.

Noise isn't bias. Bias is systematic: it errs the same way every time, like a mapping model that puts every client into too risky a portfolio. Noise errs in more mysterious – and therefore less easily manageable – ways.

Accurate



Biased



Noisy



Noisy & Biased



Adapted from Kahneman et al, HBR, 2016

Upfront assessments are necessary but insufficient, and often overplayed. Suitability reflects circumstances. And circumstances change constantly. Sometimes changes are imperceptible. Sometimes pandemics happen.

Because this is inherently complex, we are drawn towards the sanctuary of the status quo. Overemphasising initial assessments makes investment solutions over-fitted to the client's circumstances at that single point in time, and unresponsive to subsequent changes. They drift away from what is suitable over time. In times of crisis this drift can become a dash.

Risk tolerance – a client's long-term willingness to take risk – is a largely stable and readily quantified attribute. It should form the foundation of investment suitability. But suitability should also be responsive to the client's financial circumstances, which determine their risk


capacity (their financial ability to take risk) – especially during an investment journey.

But risk capacity has many moving parts. Studies on multi-attribute decision-making show that even when people think they're assimilating evidence from all sources, they're really just filtering down to the few that stand out. Those few aren't consistent over time, let alone over different decision makers.

Establishing frameworks to drive consistency in diagnosing situations doesn't mean giving every client the same answer. It means those answers need to be within boundaries defined by a clear diagnosis of the problem. There are multiple paths towards remedying any situation, depending on client personality, circumstances, and engagement.

Identifying noise isn't about eradicating inconsistencies. It's about eradicating unjustifiable ones and evidencing justifiable ones.

Noise isn't bias. Bias is systematic: it errs the same way every time, like a mapping model that puts every client into too risky a portfolio.



Survey design

Survey design

Developing hypothetical client profiles

To detect noise in investment advisers' decisions, and to understand where it comes from, we designed a roster of six imaginary clients. With brief summaries of relevant (as well as some irrelevant) information, advisers were asked to make various judgments about the clients.

We chose to create six hypothetical clients because any more would start to feel onerous for advisers, and potentially reduce survey completion rates, while any fewer would inhibit the potential for gleaning valuable insights.

This methodology allows us to test the effects of changing a single variable while keeping

everything else constant. For example, clients 1 and 4 are identical, except for their risk tolerances: the former is a 6 out of 7 ('High') and the latter is a 4 ('Medium').

Similarly, clients 2 and 5 are the same, save for the former's much more valuable property (with no mortgage outstanding); and clients 3 and 6 vary only in that the latter has a far more ambitious spending goal.

This set of three pairs enabled us to examine, in particular, the degree to which advisers were influenced by risk tolerance and risk capacity – and within risk capacity, the extent to which non-investible assets and goals affected adviser thinking.

Roster of imaginary clients

		Client 1	Client 2	Client 3	Client 4	Client 5	Client 6
Bio	Age	68	35	45	68	35	45
	Ethnicity	Black SA	White SA	Black SA	Black SA	White SA	Black SA
	Education	Doctorate	Bachelor's	Bachelor's	Doctorate	Bachelor's	Bachelor's
	Location	Jo'burg	Pretoria	Durban	Jo'burg	Pretoria	Durban
	Marital status	Married	Single	Married	Married	Single	Married
	Children	1	0	3	1	0	3
	Employment	Retired	Employee	Entrepreneur	Retired	Employee	Entrepreneur
Financial personality	Investing exp.	3	3	6	3	3	6
	Risk tolerance	6	2	7	4	2	7
Financial situation	Risk perception	6	4	2	6	4	2
	Portfolio value	8	0,5	5	8	0,5	5
	Home value	32	50	20	32	10	20
	Mortgage value	6	0	10	6	6	10
	Business value	0	0	20	0	0	20
	Income	3	2	4	3	2	4
	Expenditure	1,5	1	0,5	1,5	1	0,5
Spending goal	20	5	12	20	5	40	

All values in the table are in millions of Rands. Investing experience, Risk tolerance, and Risk perception were each on a 1–7 scale with the following definitions:

- Investing experience – level of previous investment experience: 1 = Very little; 7 = Extensive.
- Risk perception – how risky does this client perceive the market to be? 1 = Not at all risky; 7 = Very risky.
- Risk tolerance – willingness to accept the possibility of worse long-term outcomes for a greater chance of higher long-term returns: 1 = Very low; 7 = Very high.

Advisers were also shown short vignettes for each client, to provide further colour and hint at attributes such as emotional capacity to take risk. For example, client 1 was described as follows:

Tsepho is a married and retired man. He owns his own house, valued at R32M, and would like to leave an inheritance of R20M to his child.

Tsepho has been investing for many years, but is largely disengaged and passive. He perceives markets to be risky, but does not monitor his investments much until times of market stress.

In the actual implementation of the survey, advisers were shown imaginary client names corresponding to randomised genders: 50% male, 50% female. This was to detect any potential gender bias in the investment recommendations, which has been observed in recent research. Though not our primary objective in this noise audit, the ability to randomise names in this way for a subtle gender cue made detecting any bias possible.

Limitations

The drawback of using only a small number of imaginary clients is that a correspondingly small number of variables can be tested. Furthermore, when comparing adviser responses across pairs of clients, it may be difficult to discover precisely which differences are having a given effect.

We must also acknowledge that taking a survey of this type is not the same as giving investment advice in real life. Professional advisers develop deep relationships with their clients over many years, and get to know their personalities and financial circumstances in ways that cannot be replicated with a mere table of figures. It's also highly likely that they will not give the same level of care and attention to imaginary clients as their real ones.

That said, there are plenty of distractions and difficulties in the real world, and so running a few calculations for a survey could be seen as comparatively straightforward, filtering out many potential sources of noise. Though they may be imaginary clients, we should still expect the recommendations they receive to be consistent, and find a reasonable degree of consensus amongst investment advisers.

It's also highly likely that they will not give the same level of care and attention to imaginary clients as their real ones.

How noisy are investment advisers' recommendations?

How noisy are investment advisers' recommendations?

Suitable risk level


When presented with the same hypothetical client information, advisers gave remarkably different judgments on how much investment risk would be suitable.

For four of the six clients, there was at least one adviser who recommended a 'Very low' level of risk, yet another proposed 'Very high'. Indeed, for client 5, advisers were almost

evenly split between recommending a lower, medium, or higher level of risk, suggesting strong disagreement.

For other clients, there was more harmony: almost four out of five advisers judged clients 3 and 6 to be suitable for a higher-risk portfolio. This makes sense, as, unlike the others, both these clients perceive market risk to be low, and are business owners. But even here, 6-8% of advisers opted for a lower-risk solution, despite the client having 'Very high' risk tolerance.

Suitable risk level as judged by investment advisers




... one adviser recommended a 'Very low' level of risk, yet another proposed 'Very high'.

Risk capacity

Advisers were even more split when it came to considering a client's risk capacity – a crucial input into the calculation of suitable investment risk.

For example, on a scale from 1 to 7, clients 1 and 2 were given risk capacity ratings from 2 to 6 with almost equal frequency. Only for clients 3 and 6 was there any clear consensus: over half of advisers judged this pair to have 'High' or 'Very high' risk capacity.

Risk capacity as judged by investment advisers



Financial personalities

Ability and willingness to take risk are the most important factors to assess and combine when making investment recommendations. But other personality traits matter too, as they can have significant implications for both the most prudent level of risk to recommend, as well as how to guide a client towards a wise decision.


Composure – a client’s emotional capacity to handle market fluctuations – is related to risk tolerance, but is more focused on the tendency to panic in the shorter term. When asked to judge the composure of hypothetical clients 1 and 2, advisers were almost equally split across the middle three bands of a five-point scale.

Indeed, across all clients, only 4% on average judged them to have either ‘Low’ or ‘High’ composure. This suggests a reticence to reach for extremes, with no clear notion of whether some clients fall either side of the norm. This makes sense, given that advisers were given limited information relating to emotional capacity to take risk.

That said, the clearest indicator of composure amongst the information given – the client’s stated risk perception – displayed a weak correlation to adviser assessments of composure. There was, however, a stronger link between their appraisal of client composure and risk capacity, which theoretically should be unrelated since a client’s financial ability to take risk (ie their risk capacity) shouldn’t be related to their emotional responses.


... other personality traits matter too, as they can have significant implications ...

Composure as judged by investment advisers




A similar aversion to extremes is seen in adviser assessments of both the investing knowledge of each client, as well as the degree to which each would like to have control over their investments. However, in these cases, there was at least consensus as to whether each client was lower or higher than average.

Knowledge as judged by investment advisers




Desire for control as judged by investment advisers



Portfolio recommendations

Given the level of dispersion of what advisers consider suitable risk, it should come as no surprise that their high-level asset allocation proposals are similarly scattershot. Recommended equity weights for clients 1 and 3 ranged all the way from 0 to 100%, for example.

Recommended equity weight by investment advisers



This points to another layer of noise faced by advisers: even if there is agreement among advisers about how much risk is suitable, they may disagree on what kind of portfolio would represent this.

Interpretation of risk


Do advisers recommend roughly the same proportion of equity if they are agreed on the suitable risk level?

A certain amount of dispersion is rational here: advisers may have differing expectations of risk

and return for each asset class. But only to an extent: for example, although there is no agreed-upon industry definition of a moderately risky portfolio, somewhere between 40-60% equity is considered typical.

This corresponds closely to the inter-quartile range of responses amongst those advisers who rated a client's suitable risk as 4 out of 7: an equity allocation of 45-60%. However, it follows that half the advisers opted for an equity weight outside this spread, and indeed the full set of responses ranged all the way from 10% to 80%.

Recommended equity weight per suitable risk level



Overall, we see a sensibly upward-trending allocation to equity, but there is evidently a wider spread of recommendations than could reasonably be expected by differing opinions on prospective risk and return.

Summary: measuring dispersion

To better understand and compare the variation in responses that advisers give, we need some kind of measuring stick. We can treat the percentage of responses that advisers gave for each level of a scale as if they were weights in a portfolio.

By doing so, we can use measures such as entropy. Like the Herfindahl-Hirschman index, it is used to measure diversification (or the lack of it). This makes it well-suited to checking whether responses are uniform or spread out.

We can normalise entropy to a scale from 0 (everyone gives identical responses) to 1 (the responses are completely evenly distributed across all the options). This gives a much more immediate sense of how noisy responses are, and allows us to compare noise levels across the

various client characteristics that advisers were asked to judge.

Entropy: example

To get a better feel for this measure, let's focus on the suitable risk level judgments for client 1. Advisers are mostly split across four responses, from 'Low' to 'Medium-High' (with fewer advisers within this range recommending a 'Low' level of risk). Only a small number of advisers opted for 'Very low' or 'High' risk.

Suitable risk level advocated for client 1



On a scale of 1 to 7, where 1 is everyone picking the same risk level, and 7 is exactly one-seventh of advisers picking each option, the entropy score was 4,8 - ie closer to totally random than totally consistent.

Suitable risk level for client 1: raw entropy




To give another example, if instead the adviser responses were equally split 25% in each of the four categories 'Low', 'Medium-Low', 'Medium' and 'Medium-High', they would have scored exactly 4.

We can convert this number to a normalised scale from 0 (no noise) to 1 (totally noise), to make it easy to compare across variables with differing numbers of categories. Thus in this case, we arrive at a normalised entropy score of 0,63, which can be interpreted as 'fairly noisy'. Indeed, anything above 0,5 would probably be considered undesirably high.


Entropy: analysis

We can see from the entropy measure that adviser responses are fairly noisy across the board, especially when asked to judge risk capacity. Recommended risk levels are less noisy, but then more noise is introduced by the translation into an equity allocation for the client's portfolio, as we saw earlier.

Entropy of adviser responses (higher is noisier)



Entropy of adviser responses, averaged across clients (higher is noisier)





What influences advice?

What influences advice?

Ultimately, clients paying for investment advice receive a portfolio recommendation. This is the culmination of a series of assessments and adviser judgments – risk tolerance and capacity, financial personality, and the suitable risk level for their investments.

To understand the noise of portfolio recommendations, we must therefore break down the advice journey and analyse what factors play an influential role at each stage.

As we do so, we can consider each stage through two lenses: observable factors and unexplained variation.

First: how much of the adviser's judgment is driven by observable factors (differences in either the client case study, or the advisers themselves)?

This includes both factors that should and shouldn't influence the suitable solution.

Observable factors that should influence the suitable solution include, for example, the risk tolerance of the client.

Because there may be reasonable disagreement about the weight that certain characteristics should have in determining the right answer, it is difficult to be precise about the degree to which any input should change the recommended investment solution. But we can differentiate between:

- factors that are influencing the solution in the right direction (eg if clients with higher risk tolerance are, all else constant, given higher risk);
- factors that should influence the right answer but don't appear to be doing so (for example, if clients with different levels of risk tolerance didn't seem to be getting different risk recommendations); and

- factors that seem to be used wrongly in arriving at a recommendation (for example, if clients with higher risk tolerance are getting on average lower risk recommendations).

Observable factors that shouldn't influence the suitable solution could be either about the client (eg the gender of the client shouldn't matter for their suitable risk level), or the adviser (eg perhaps older or younger advisers tend to give clients different solutions). Of course, the suitable investment solution for any client should only be a reflection of their needs, never the characteristics of the adviser.

Second: how much appears to be down to unexplained variation?


Here too we can discern two types of variation in advice: idiosyncratic divergence across advisers, and inconsistencies within each individual adviser.

Advisers may be different from each other in ways we can't observe in the data, and these differences could lead different advisers to (wrongly) give different solutions to their clients.

To understand the noise of portfolio recommendations, we must therefore break down the advice journey ...

Each individual adviser could give inconsistent solutions to clients based not on the client's circumstances, but the adviser's. For example, imagine an adviser whose mood was heavily influenced by the weather, and which, unknown to themselves, led to a tendency to give the same client different answers depending on whether they met on a sunny or rainy day.

Of course, we can't accurately assess the magnitude of all these effects using the data from the survey. But we can provide indications of where advice seems to be more or less consistent, and some idea of where any inconsistency comes from.



Predictive factors: signals along the advice journey

To determine how the variability in advice is driven by observable factors, we can look at all the data collected in our survey. We use statistical techniques (in this case regression analysis) to determine to what degree differences in the assessed suitable risk level is driven by each of the input variables.


The data we have can be grouped into four main categories:

1. **Client characteristics** - the variables that survey respondents were given for each client. For example, the client's gender, age, balance sheet, risk tolerance, etc. Some of these should influence the suitable investment solution in a particular direction, while some definitely shouldn't.

- 2. **Adviser characteristics** - what we know about each adviser as an individual, including demographic variables (education, etc), industry variables (compensation structure, firm type, client segment, years of experience), and psychometric variables (six aspects of financial personality, as well as their self-judgment of their knowledge).
- 3. **Adviser judgments** - the subjective scores respondents gave to each client case study, assessing them on their composure, desire for control, and knowledge.
- 4. **Adviser recommendations** - for each client's risk capacity, suitable risk level, and suggested asset allocation, as well as the adviser's confidence that their recommended solutions were correct for each case.

In the chart below you can see how much different sets of data explain the variability of the assessed suitable risk levels.

Factors influencing suitable risk level assessment



If advice across the whole set of responses was not noisy, we should expect to see that different suitable risk levels for each client should almost exclusively be explicable by differences in the information given to advisers on each client - the client characteristics - with little influence from adviser characteristics.

Instead, we see that:

- 1. By itself, the information most relevant to determining the right answer explains only 28% of the variability in adviser recommendations.
- 2. Even considering all data, including variables that shouldn't have any influence on the suitable risk level, the explained variance is only 56%. There is a large amount of variation in the recommended risk level that doesn't seem to be driven by the information advisers were given. That is, there seems to be a large amount of subjectivity and noise in the recommended risk level.

However, this is by no means all worrying news for advisers.

Quite a large chunk of explanatory power is made up by adviser judgments of the client's situation. We can see that including the adviser's judgments of the client's composure and knowledge raises the explained variability to 39%. And when the advisers' assessments of the clients' risk capacity are included, the proportion explained rises to 54%.

... there seems to be a large amount of subjectivity and noise in the recommended risk level.

¹Measured by the R² of the regressions

This implies that judged risk capacity (the assessment that should probably have the greatest effect on the recommended risk level) on its own increases the explained variability substantially.

And although the recommended risk level is not being driven in a very clear and direct way from the client characteristics, advisers are nonetheless assimilating this information (possibly in quite individual and idiosyncratic ways) into higher level perceptions of each client's overall situation and these assessments are driving the risk recommendation.

Also important is that, while including information on advisers rather than clients does affect the recommended risk level (which it shouldn't), this only explains 2-3% of the variation.

Finally, for each category of variables we can ask: which specific items of data are having an effect? And, importantly, are these effects sensible, or influencing recommendations in the wrong way?

Client characteristics

Although the raw client characteristics only explain 28% of variability in the recommended risk levels, it is at least encouraging that advisers are very significantly responsive to differences between clients. Are these responses in a sensible direction?

The answer is mostly, yes. Client characteristics do seem to be on the whole influencing recommended risk in the right way, though perhaps not with the magnitude we would expect.

For example, we see that clients with higher risk tolerance do get given higher risk levels, but this effect comes entirely through risk tolerance affecting the advisers' judgement of risk capacity. In other words, the effect is indirect: increasing risk tolerance by one category on a 7-point scale leads advisers to an assessment of risk capacity that is on average 0,3 points higher, and this in turn leads to higher risk recommendations.

This might be a reasonable outcome, but does indicate that advisers don't on the whole have

a clear model of how risk tolerance and risk capacity are distinct from each other, and how each should independently influence the recommended risk level.

To put it another way, two hypothetical clients with the same financial circumstances should have the same risk capacity (financial ability to take risk), but could still have different levels of willingness to take risk (risk tolerance) and thus should be recommended a different risk level.

Our data suggest that risk tolerance and risk capacity are not being independently used in determining the risk recommendation but conflated together.

In addition, the final effect of higher risk tolerance on the recommended risk level is significantly weaker than might be expected, given the important role risk tolerance should play.

Having large home equity (such as the difference between client 2 and client 5 in the case studies) is another variable that affects the recommended risk level in the right direction (higher home equity leads to greater recommended risk) via the adviser's risk capacity assessment. This is exactly what we should expect.

Another variable that influences recommended risk levels, but only via the assessment of risk capacity, is more concerning: client case studies with female names were assessed as having slightly higher risk capacity, leading to potentially different advice based only on client gender.

Adviser judgement

We have already seen that the advisers' assessment of each client's risk capacity is a strong influence on the risk level they recommend. And also that some client characteristics are influencing this risk level only via their effect on assessed risk capacity.

This is consistent with what we'd hope to see: risk capacity should be a primary driver of recommended risk levels. Quantitatively, for

every point higher the assessed risk capacity, the recommended risk level is 0,36 points higher (both on a 7-point scale).

But advisers' assessments of other aspects of each client are also influential. Clients who are perceived as being more composed are recommended higher risk levels. And even more strongly, clients who are assessed as being more knowledgeable are recommended more risk.

It is quite reasonable that both of these affect the suitable risk level, but in both cases the extent to which this seems to be the case is surprisingly large, given that the correct answer should be much more about the client's long-term willingness and ability to take risk, rather than their knowledge (the delegation of which is part of the reason for seeking advice) or their short-term emotional ability to take risk (which is much more effectively and cheaply dealt with through good client communication and engagement, rather than reducing the long-term portfolio risk).

Adviser characteristics

The characteristics of the adviser really shouldn't drive the suitable solution for each, and indeed differences between advisers drive only a small portion of the variation in recommended risk levels.

Nonetheless, there are certain characteristics of advisers that do seem predictive of the advice given:

- Married advisers recommend slightly lower risk levels than advisers who are single;
- University-educated advisers have lower risk capacity assessments on average;
- Salaried advisers give higher recommended risk levels than those on commission or fee-based.

However, interestingly, how experienced the adviser is, or how many clients they serve seems to make no significant difference to the advice delivered.

¹See *New Vistas in Risk Profiling* (Global CFA Research Institute) for a technical discussion of the differences between risk tolerance and risk capacity and how they should be independently used in arriving at the suitable risk level.

Lastly, we can examine whether the financial personality of the adviser leads them to make different risk recommendations. Of the dimensions of financial personality we measured, only risk tolerance led to a reliably significant difference in the recommended risk levels: advisers who themselves are more tolerant to risk tend to pass it on to their clients.

What explains differences in asset allocation?

We have so far examined the variables that help explain why different clients get recommended different risk levels. However, in the advice process there is also the secondary stage of implementing a given risk level into a portfolio. Do different advisers build different asset allocations for a given client risk profile?

Allocation to equity

As expected, once we know what suitable risk level the adviser has arrived for each client, we can explain a great deal of the variation between the suggested equity allocations. And reassuringly, the risk level is very closely related to the asset allocation (with a correlation of 78%).

But this does still leave some differences in how advisers construct portfolios for a given risk level.

Other variables that affect this equity allocation include:


- Differences between clients still influence the equity allocation, even for a given risk level. In other words, advisers often build different portfolios for two clients, even when they have been given the same risk level. In particular:
 - Clients with the same risk profile but with higher risk capacity are given higher equity allocations;
 - Clients with higher perceptions of risk (ie the client thinks markets are more risky) are given lower equity shares than clients with the same assessed risk profile.

- The financial personality of the adviser affects the equity allocation they recommend:
 - Interestingly, more risk tolerant advisers recommend lower equity allocations for given risk profiles (though, as we've already seen, they tend to arrive at a higher risk profile). This could be because of some concern for not being seen to push clients into more equity than they're comfortable with. It is perhaps easier to justify nudging clients into

a higher profile than it is to be seen given too much equity for a given profile.

- However, advisers who have higher composure do recommend significantly more equity for each risk level. This makes a lot of sense as these advisers are likely to be much less anxious about short-term volatility and more focussed on long-term risk vs return.

Cash weight recommended by investment advisers



Allocation to cash

As might be expected, the suggested cash allocation is in many ways driven by the same variables as the equity allocation, but in the opposite direction. For example, a higher risk level leads, quite reasonably, to a lower cash allocation.

Advisers who are single tend to recommend more cash.

- Unlike the equity allocation, the cash allocation does not vary with the adviser's risk tolerance (but does increase significantly for lower composure advisers).

However, in a few ways the allocation to cash has distinct drivers:

- It is less sensitive to client information, and less explained by observable data – advisers perhaps have their own ideal cash levels that are relatively independent of client characteristics.

Advisers who are single tend to recommend more cash.

Allocation to South African vs global equities

The proportion of equities that advisers would put into South African equities ranges all the way from nothing to everything for all six clients. There is, however, a substantial spike at the 50/50 level, and allocations of greater the 50% to South Africa are more common than allocations of less than 50%.


Only 15% of the variation in this allocation can be explained from the observable data, leading to the conclusion that adviser opinions on this matter are highly idiosyncratic and individual.

However, we can identify a few variables that affect this allocation:

- Clients with higher risk levels are recommended lower allocations to SA.

- Fee-based advisers on average recommend less SA equity.
- Advisers with more clients recommend slightly more SA equity.
- Advisers with lower composure scores recommend less SA equity.
- More experienced advisers recommend less SA equity; but on the other hand those for whom it is a long time since their last professional designation was awarded tend to recommend more.
- And, perhaps most interestingly, those advisers who have a strong internal locus of control (ie, they have a conviction that success is due to hard work and ability, not luck) recommend significantly less SA equity.

SA equity proportion recommended by investment advisers



Inexplicable variation: noisy advice

Unidentified idiosyncrasies

We have identified several factors that correlate with investment advice – some sensibly and intuitively, others less so. However, there remains a great deal of variation in the responses from one adviser to the next.

At least part of this variation can be attributed to differences between advisers that our analysis did not detect – for example, personality traits beyond those related to investing. Such analysis is beyond the scope of this noise audit.

However, it may well be the case that there are patterns of responses common to certain sub-groups of advisers that are obscured when attempting to identify broad trends across the entire surveyed sample.

For an analysis of these clusters, please see the 'Adviser archetypes' section.

Inconsistencies

One type of noise is especially troubling: inconsistency. It is one thing for there to be unpredictable variation across advisers, it's quite another for the same adviser to give answers that vary from one moment to the next.

To examine whether advisers give consistent answers when faced with the same information,


we can exploit the design of the hypothetical clients - which consist of pairs that differ only in a single variable. We can then check to see if advisers factor in this change in a sensible manner.

Consideration of risk tolerance

Hypothetical clients 1 and 4 are identical, except for their risk tolerances: the former was described as a 6 out of 7 ('High'), the latter a 4 ('Medium'). So, we should expect client 4 to be uniformly suited to a lower level of investment risk - as no matter how advisers integrate risk tolerance into their decision, there's only one rational outcome.

Perhaps surprisingly, almost one in five advisers decided that client 4 should have an investment portfolio riskier than client 1 - despite the only

Adviser suitable risk level responses: client 4 vs client 1
Client 4 is identical save for lower risk tolerance



difference (lower risk tolerance) suggesting the opposite. A further one in three advisers gave them the same suitable risk levels, which suggests at the very least an under-weighting of risk tolerance in the decision.


Consideration of non-investible assets

Clients 2 and 5 differ only in respect of their home value and mortgage outstanding. Client 2

has an unencumbered home worth R50 million, whereas client 5's is worth R10 million with a mortgage of R6 million.

We observe that roughly one in five advisers appraise client 5's risk capacity as being lower.

Adviser risk capacity responses: client 5 vs client 2
Client 5 is identical save for less home equity



The precise effect that non-investible assets should have on one's risk capacity - and hence suitable risk level - is debatable. Owning a property with a large net value clearly increases one's capacity relative to the same client who does not have this asset, though advisers might reasonably differ as to how much this should affect the risk capacity assessment.


However, there can be no argument that, if an adviser considers one client to have higher risk

capacity, they should also recommend they take on more risk - because all else is equal.

We find that in fact, 6% of advisers gave incompatible responses: they consider client 5 to have more/less risk capacity, but assign them a lower/higher suitable risk level.

A further two in five advisers judged them to have the same risk capacity but differing suitable risk levels, or vice-versa.

Incompatible adviser risk capacity/suitable risk level responses: client 5 vs client 2
Client 5 is identical save for less home equity




Adviser archetypes

Adviser archetypes

We have so far broken down adviser decisions into predictable and unpredictable influences – but only when considering their responses in aggregate. It may be that there are various types of adviser, behaving in particular ways that make them identifiable as a subgroup, but who may be drowned out in the crowd.

Statistical techniques known as ‘clustering analysis’ can be used to tease out such cliques. There appear to be four distinct clusters of advisers, whose judgments on suitable risk levels are more like those within their group than the others:



The following chart shows how much each cluster differs from the overall survey average with their recommended suitable risk levels. It shows the number of standard deviations away from the mean, to help illustrate how close to the typical response (or unusual) each cluster is.

Suitable risk levels per cluster for each client (Standard deviations from average)



Cautious

This cluster of advisers gave lower suitable risk recommendations to almost all clients, except for clients 3 and 6 (where the 'Unsure' advisers may have distorted the picture). This is quite possibly related to the fact that this group has lower risk tolerance and composure themselves, so they may well be projecting their own characteristics onto their clients.

These advisers seemed to give especially low risk recommendations to clients 1 and 4. It is hard to discern precisely what is causing this, though we might speculate that it is the relatively old age of these hypothetical clients that causes concern.

Unsure

These advisers give judgments that change little from client to client, and much lower risk recommendations to clients 3 and 6.

We can only speculate as to what exactly it is that these 'Unsure' advisers are picking up on for this pair of clients. But we note that they give them much lower risk capacity scores than do their peers. So, it must be something to do with risk capacity, yet it can't be the ambitious spending goal of client 6, because client 3 does not share it.


Given the descriptions of these two clients, the most plausible candidates are the number of children they have (three, compared to the other clients who have one or none), or the sizable business asset that the client owns.

Furthermore, we note that this is the only cluster of advisers to assign client 4 a higher suitable risk level on average than client 1, despite them being the same except for the former having a lower risk tolerance. This suggests a degree of inconsistency not seen amongst the other archetypes.

... they may well be projecting their own characteristics onto their clients.


After accounting for factors that consistently explain adviser recommendations - in one way or another - this Unsure group of advisers give the most varied responses (as reflected in the 'residual' difference between what one would expect from those factors and their actual answers).

Noise of each cluster Average sum of square residuals per cluster*



However, interestingly, the 'Unsure' advisers are the group giving the most consistent responses in absolute terms:

Average deviation from mean suitable risk level



This gives us a fascinating insight: these advisers are effectively the most 'noisy' not because their appraisals are all over the place – far from it, they are remarkably stable in their opinions – it's because the clients are different, and require different solutions.


In other words, these advisers are like archers firing off arrows in roughly the same direction every time – at a moving target. Hesitant as to whether they'll hit the bullseye, they avoid recommending especially high or low risk.

Risk-tolerance focused

These advisers are more responsive than most to the risk tolerance of the clients: their suitable risk recommendations appear to be more influenced by willingness to take risk than the other clusters, after accounting for judged risk capacity.

As a result, their recommended suitable risk levels came out closer, on average, to the risk tolerance described for each client.

Average difference between suitable risk levels assigned per cluster and client risk tolerance



As such, these advisers were the only ones to give suitable risk ratings a significant amount both above and below average across the clients. Significant, but not extreme.

As for risk capacity: these advisers' opinions deviated the least from the average, suggesting a generally orthodox approach.

Relaxed

The 'Relaxed' cohort give much higher risk recommendations to clients 2 and 5, who differ only in that the former has a large, unencumbered personal property. These advisers do consider both these clients to have greater risk capacity than do their peers, but not enough to fully account for the gap.

It could well be that these advisers noted that both clients 2 and 5 are described as relatively young, single, and with no dependents. Or perhaps they gave relatively less weight to the low risk tolerance of this pair of imaginary clients, given that these advisers are sanguine about market risk themselves.

It is also worth observing that the 'Relaxed' advisers are rather noisy in their responses – while not quite as much as the 'Unsure' cluster, they are still somewhat unpredictable.

Conclusion and recommendations

Conclusion and recommendations

Noise vs bias

Advice appears to be predictably influenced by a broad range of factors – both relevant and irrelevant. When incorporated incorrectly or inappropriately, this is bias.

Beyond this, there is also noise borne of random error. But an unbiased lack of accuracy can also stem from giving consistent answers to changing problems. In the context of this survey, it arose from missed opportunities to refine advice based on relevant information provided.

We summarise our main findings below:

	Relevant information	Irrelevant information
Bias (approx. 20% of variation)	Incorrectly incorporated For one in five advisers, risk tolerance affected the recommendation in the wrong direction.	Partiality Adviser financial personalities projected onto their clients – eg higher adviser risk tolerance passed on to client.
Noise (approx. 50% of variation)	Missed opportunities “Unsure” cohort of advisers giving similar answers to different clients.	Random error Advice shows no clear model of how to integrate risk tolerance and capacity.

Possible explanations for inconsistencies

There are all sorts of well-researched reasons why humans sometimes make poor decisions. We are not robots – our brains are deeply connected to the world around us.

The environment we’re in, the time since our last meal, our mood – there are a great many subtle influences at play which can lead to seemingly random variations in our decisions.

It is possible that at least some advisers ‘got used to’ the survey as they progressed, and refined their approach. Furthermore, we cannot exclude the possibility that some advisers were influenced by the most recent judgments they had made – for example, by giving a client a higher risk level than they otherwise would, after having recommended increasing levels of risk for the preceding clients.

However, there are always framing effects in the real world, too. And there are simple, practical steps one can take to improve decision-making performance, whatever the situation.

How to cancel noise

Identifying inexplicable variation in financial advice is one thing. Doing something about it is

another. Unlike biases, one cannot so easily self-correct noise through self-awareness.

The most effective noise-cancelling remedy is to employ algorithms where there are decisions to be made. By rigorously following a set of pre-defined rules, advice can be made much more consistent.


Fortunately, the field of investment advice is well-suited to the deployment of algorithmic tools to aid decision-making. Once a specific framework for the measurement and synthesis of risk tolerance, risk capacity and other relevant factors has been developed, it can be run at both speed and scale.

This emphatically does not mean the end of the human adviser – clients place immense value on their relationships, and trust is an invaluable commodity no robot can replace.

Ultimately, the role of algorithm is to function as a kind of decision prosthetic. Like the Decision Review System (DRS) used in cricket, or the Television Match Official (TMO) in rugby, technology can be employed to greatly increase consistency and accuracy. But in the end, when the margins are extremely tight, it’s the umpire’s call.

So should it be in the world of investment advice.

³Noise: How to Overcome the High, Hidden cost of Inconsistent Decision Making (hbr.org)



momentum

investments

Special thanks to Nedbank Financial Planning for their support in driving critical responses for this landmark study.



Disclaimer: The information used to prepare this document includes information from third-party sources and is for information purposes only. Although reasonable steps have been taken to ensure the validity and accuracy of the information contained herein, Momentum Metropolitan Life Limited does not guarantee the accuracy, content, completeness, legality or reliability of the information contained herein and no warranties and/or representations of any kind, expressed or implied, are given to the nature, standard, accuracy or otherwise of the information provided.

Neither Momentum Metropolitan Life Limited, its affiliates, directors, officers, employees, representatives or agents (the Momentum Parties) have any liability to any persons or entities receiving the information made available herein for any claim, damages, loss or expense, including, without limitation, any direct, indirect, special, incidental, punitive or consequential cost, loss or damages, whether in contract or in delict, arising out of or in connection with information made available herein and you agree to indemnify the Momentum Parties accordingly. For further information, please visit us at momentum.co.za. Momentum Investments is part of Momentum Metropolitan Life Limited, an authorised financial services and registered credit provider, and rated B-BBEE level 1.