Research
Foundation
Briefs

NEW VISTAS IN RISK PROFILING

Greg B. Davies

**CFA Institute
Research
Foundation**

For more on upcoming Research Foundation publications and webcasts, please visit www.cfainstitute.org/learning/foundation.

Research Foundation monographs are online at www.cfapubs.org.

NEW VISTAS IN RISK PROFILING

Greg B. Davies

CFA Institute
Research
Foundation

Statement of Purpose

The CFA Institute Research Foundation is a not-for-profit organization established to promote the development and dissemination of relevant research for investment practitioners worldwide.

CONTENTS

NEW VISTAS IN RISK PROFILING

Greg B. Davies
Centapse

A CONFUSED LANDSCAPE

Risk profiling is one of the most fundamental aspects of determining a suitable investment solution for an individual. It is also one of the most misunderstood.

This situation results in many ill-advised approaches, both to determining various components of a client's risk profile and to using them in arriving at an appropriate solution.

Among these ill-advised approaches are

- using a client's actual behaviour (their "revealed preferences") to determine the *risk tolerance* used to construct their long-term portfolio;

- using components of a client's attitudes toward risk other than long-term risk tolerance as inputs into portfolio "optimisation," rather than helping the client to mitigate and control them;

- trying to elicit risk tolerance on subcomponents of a client's overall wealth, rather than as a feature of the person as a whole;

- using over-engineered and unstable approaches to establishing risk tolerance when simpler measures would be not only sufficient but also better;

- putting too much effort into spuriously precise measures of risk tolerance and far too little into understanding *risk capacity*, which in many cases is the more important component of a *risk profile*; and

- using "required" returns as inputs to determine the investment solution, rather than guiding the client to reasonable expectations as an *outcome* of the client's circumstances and of the returns available in markets.

Much of the problem lies in the rather imprecise and ambiguous way the industry uses the terminology around risk profiling, with terms such as *risk tolerance, risk capacity, risk appetite*, and *risk attitude* applied to a range of different concepts in different parts

of the world, in different organisations, and, indeed, sometimes within an organisation in different conversations with a single client.

However, even where terminology is stable, there are areas of confusion. For example, the industry lacks a coherent model of risk capacity and how it relates to risk tolerance (both as defined below). What gets measured gets managed, so risk tolerance is frequently called on to do much more work than it should in determining a suitable portfolio solution. Though a full model of risk capacity is beyond the scope of this article, it does have implications for how we think about risk tolerance and how we go about measuring it.

Also, there is a frequent tendency to use short-term behavioural responses in "optimising" a portfolio—that is, building a portfolio solution around the assumption that an investor's observed behaviour is somehow "desirable" and should be used as an input to the optimization process. Instead, one should try to make clients more comfortable with a long-term approach that treats such responses as undesirable. For example, attempts to use *loss aversion* or *ambiguity aversion* as inputs into an optimal portfolio solution often fail to recognize that there are aspects of risk attitudes that are helpful in determining the "right" investment and that there are aspects that are harmful. We should help investors to *overcome* the harmful ones to prevent their behavioural responses from getting in the way of good long-term outcomes. These aspects of the client's overall risk profile should refine, not define, the investment solution.[1]

In the first part of this article, I shall discuss how we should think about the crucial elements of (a) risk tolerance, (b) behavioural risk attitudes, and (c) risk capacity. In the second, I use a simple thought experiment to examine a stripped-down investor situation and define the essential features and exact role of each of the components of an investor's overall risk profile. In doing so, I shall use terminology that is broadly consistent with that used by regulators to stipulate how each concept should be used in delivering a *suitable* investment solution. For example, risk tolerance will refer to that subcomponent of risk attitudes that should be used in determining the long-term "right" investment solution, rather than to any attitude toward risk that the investor happens to have over the investment journey. Others may choose to apply different terms to concepts unearthed in the thought experiment; sadly, I suspect a globally consistent terminology might be too much to hope for. But if we can at least reach agreement on what the essential components are, however named, then we will have made progress.

Armed with a clear distinction of what we mean by the various elements of risk profiling, it becomes a much simpler task to ask how we might best go about eliciting these

---

[1]"Beyond Risk Profiling" (2015).

elements. The final part of the article will examine options for eliciting and measuring risk tolerance and will consider some promising avenues for future methods.

REVIEWING THE LANDSCAPE

To undertake risk profiling effectively, we need to have a clear idea of what it is we're trying to understand. Broadly speaking, a comprehensive risk profile requires understanding two characteristics of a client:

1. their willingness to take risk

2. their ability to take risk[2]

But what exactly do we mean by these terms, and what do we need to understand about the client in order to put them to practical use?

WILLINGNESS TO TAKE RISK

The first aspect, willingness to take risk, is often variously, and ambiguously, referred to as risk attitude, risk tolerance, risk aversion, or risk appetite. In a sense, it doesn't matter too much what we call it. What is important is that there is a crucial distinction to be made between

a. *risk tolerance* (an investor's stable, reasoned willingness to take risk in the long term) and

b. *behavioural risk attitudes* (the unstable, behavioural, short-term willingness to take risk exhibited through an investor's actions—i.e., the investor's "revealed" preferences).

From the perspective of investment suitability, the former reflects the normative level of risk we should seek to deliver for a client over the long term.

---

[2]Required or desired returns typically should *not* form part of an investor's risk profile. In the vast majority of cases, these should be determined by whatever expected returns might be attained through efficient use of the client's appropriate risk level, which is in turn determined by his risk capacity and risk tolerance. For most of the population, the appropriate level is the highest amount of risk that can be taken given both willingness and capacity to take risk; one of the two will be a biting constraint. If a client's goals are too expensive for either her willingness or capacity to take risk, then the goals should be adjusted, not the risk profile. Increasing risk to attain "required" returns inappropriately increases the chances of catastrophic failure to reach all goals. Only where both capacity and willingness are high relative to the future needs and goals of the investor (usually only the case for particularly wealthy individuals) will there be sufficient slack in the system for required returns to be useful. Even then, there is an argument for still basing the risk level on risk tolerance, since this will provide the best security against future changes in the investor's goals and preferences.

The latter, in contrast, frequently reflects transient and context-dependent preferences[3] that would result in poor satisfaction of the investor's long-term needs if we mistakenly take them to be a guide to long-term risk preferences.[4]

This second set may contain a whole range of different influences on risk attitude, including loss aversion, ambiguity aversion, risk perceptions, and probability distortion. Apart from being unstable, these are also context dependent, myopic, present biased, and procyclical and often lead to absurd levels of risk aversion when aggregated. They are not attitudes we should want in the driver's seat for our long-term portfolio optimisation. If used to determine the "suitable" risk level of the portfolio, these attitudes can deliver an undesirable long-term solution based on an assessment of fleeting point-in-time preferences, with on average a much lower level of risk than the investor's own long-term willingness to take risk and with the added danger of buying low and selling high.

In the end, the role of suitability is to steer investors to better outcomes, not replicate (and optimise for) all the silly things they do already.

ABILITY TO TAKE RISK

In addition to ascertaining how much risk investors are willing to take, it is important to understand how much risk they are *able* to take. This concept also goes by many loosely defined labels, such as risk capacity, capacity for loss, and ability to bear loss. At their core, these are all variations on the underlying notion that the amount of risk an investor might be willing to take could nonetheless be ill-advised if it endangers the investor's ability to fund future financial commitments from the portfolio. I shall refer to this notion as *risk capacity* throughout this article, using it to encompass all the terms listed above.

Since risk capacity concerns the investor's ability to meet any future liabilities,[5] it depends on a number of elements of the investor's holistic financial position. These all help determine the degree to which the investor will be able to continue to fund future required cash flows in the event his portfolio value declines:

- First, the size of the total investable assets relative to total net wealth. Investable assets are where the investment risk resides, so the fewer investable assets clients have relative to other assets, the higher their capacity to take risk. Investors who invest only a small portion of their net wealth have more capacity to take risk

---

[3]Payne, Bettman, and Schkade (1999); Slovic (1995).
[4]Odean (1999); Davies and de Servigny (2012), chaps. 2 and 5.
[5]Apart from some infrequent circumstances where a future commitment can only be funded from a specific pool of assets.

without jeopardising future commitments. The value of these other assets may vary substantially in value and indeed is only likely to be realised when the investor's circumstances are pressured. This implies that a very conservative assessment of the value of non-investment wealth is appropriate. Nonetheless, an investor with other wealth to fall back on will have greater risk capacity than one who doesn't.

- Second, the degree to which the investor is willing or able to realise the value in non-investment assets to cover liabilities in the event that they can't be funded from the investment portfolio. More liquid assets, and those practically and emotionally easier to sell if circumstances require doing so, will provide more risk capacity than those the investor can't or won't realise the value of until they are in dire straits.

- Third, the degree to which future spending might be met using the client's expected future income,[6] which may be thought of as *human capital* or *net future wealth*. It depends on the following:

  - Future income expectations: More certain and stable future income will more reliably increase risk capacity than variable or uncertain income streams (e.g., don't count too much on distant inheritance expectations when determining today's risk capacity).

  - The future liabilities themselves: The greater an investor's goals and spending needs, the more her investment portfolio needs to provide for and, therefore, the lower her capacity to place the capital value at risk. Note that in this sense, the greater the wealth "required" to fund future liabilities, the *lower* the capacity to take risk, contrary to many approaches that use "required returns" as an excuse to *increase* portfolio risk.

  - Flexibility regarding future liabilities and expenses: If the investor has discretion over whether to incur a future expense, then expenses can be reduced or delayed in the event of poor market performance. Such an investor thus has greater risk capacity.

These last elements show the essential, and close, connection between any reasonable assessment of risk capacity and both (a) a truly holistic view of the client's current circumstances and (b) a coherent approach to goal-based investing. Risk capacity, correctly conceived, is the vital pivot between goal-based cash-flow-planning approaches and investment solutions. It turns the information gleaned from a thorough planning process into a measure that helps define the appropriate level of risk for the investment process.

---

[6]And the degree to which unexpected future liabilities are insured against.

Interestingly, this observation also reveals that even risk capacity, usually thought of as being about the investor's objective financial circumstances, is strongly behavioural. One of the most effective routes to increasing risk capacity is to manage client aspirations, expectations, goals, and future financial commitments. The less you want, the more you can take risk to grow your wealth and the more you may wind up with.

As mentioned, this article is not the place to develop a full framework for measuring risk capacity. However, understanding what it is—and isn't—is vital to being able to approach a client's willingness to take risk and to understand some of the future directions for assessing risk tolerance. It also helps us understand why developing ever more precise measures of risk tolerance may not be the best use of resources.

For most investors, risk capacity is overwhelmingly more important than risk tolerance; the right level of risk for their investments is far more likely to be constrained by their lack of capacity to cope with capital losses than by their psychological aversion to long-term risk.[7] And yet, because it is inherently more difficult to measure, with more moving parts, risk capacity is also largely neglected.

The relative importance of getting risk capacity and risk tolerance right is related to the ratio of the *flow* of income and expenditure to the *stock* of wealth. Investors with small flows in and out of their balance sheet relative to the size of the balance sheet (typically very wealthy people with large net asset values) will be likely to have high risk capacity, since they have significant wealth with which to fund future liabilities as they arise. For them, risk tolerance is far more likely to be the binding constraint.

For most investors, however, the future expenditures they will have to fund from their investible assets are substantial relative to their stock of wealth. For them, risk capacity is a far more important notion than risk tolerance, and yet it is one that many investment suitability processes largely gloss over. Even investors with high willingness to take risk should avoid doing so if the potential losses will mean they're unlikely to be able to fund future commitments. And yet the bulk of the industry debate (to which, admittedly, this article is contributing) is devoted to how to better measure risk tolerance, a stable feature of someone's personality easily measured to an adequate degree of accuracy with a simple psychometric scale.

However, part of the reason there is little agreement on how to categorise risk capacity and the two broad categories of risk attitudes (risk tolerance and behavioural responses) is that all of this appears hugely complicated. It is very easy for confusion to arise or for those with vested interests in particular views to obfuscate the debate. So, let us try to

---

[7]Short-term risk preferences are another matter, but they often lead to the investor not taking enough risk and making mistakes along the journey, rather than being a reason not to take risk.

isolate each of these three components with a thought experiment that strips away all the extraneous detail and allows us to focus on the essential features.

A RISK-PROFILING THOUGHT EXPERIMENT

Let us first imagine the simplest possible hypothetical investment situation for a single investor, Christy, and ask what the notions of risk tolerance and risk capacity mean in this context. Then we will gradually add features to the example to see what insights we can glean.

CASE 1: ISOLATING RISK TOLERANCE

Let's assume that Christy has a single portfolio of investable assets (for the simplicity of a round number, let's assume $1 million) that she wishes to invest for a fixed period of 10 years.

She has no need to draw income from the portfolio: All her future spending needs are exactly covered by her future income, and she is fully insured against any unexpected financial needs. She also has no other assets or liabilities. After 10 years she will collect the investment and move on to invest elsewhere. She will have no need at that point to withdraw or spend any part of the portfolio.

These assumptions remove, for now, considerations of risk capacity. Christy is investing the totality of her wealth, current and future, and there are no possible future financial commitments that would be jeopardised by losses, either through the 10-year period or at the end.

While this is a highly stylized and unrealistic scenario, it helps strip a complex problem back to its essentials to understand what really matters. Also, this is not much more than advisers often know about a client who brings them a portfolio to invest and doesn't wish to disclose much information about their broader financial situation or future goals. The bulk of the work in determining what is suitable has to be done by a risk-profiling process.

In addition, imagine at this point, for the sake of simplicity, that Christy will act completely on your recommendation for how the money should be invested and will pay no attention whatsoever to the portfolio in the intervening 10 years. These last assumptions are crucial for eliminating the possibility of behavioural responses to initiating the investment or to the investment journey.

What is the right level of risk for Christy's portfolio?

This bare-bones investment scenario is designed to isolate the role of risk tolerance. Risk capacity is neutralised, and emotional or behavioural aspects of risk attitude are neutralised. All that matters is Christy's willingness (now) to trade off the future chance of good financial outcomes for the possibility of bad outcomes. This is the type of problem envisioned in Markowitz's mean–variance analysis—finding an optimal trade-off between risk and return to suit the long-term risk tolerance of the investor.

Because only long-term outcomes matter, any measure of risk tolerance that is unstable, dependent on Christy's current context or emotional state, or influenced by the current state of the market is going to result in a portfolio aiming for the wrong target. And in this scenario, Christy won't even notice valuations along the way, so trying to incorporate any behavioural features that depend on her current reference point (e.g., loss aversion) will be inappropriate. What *is* required in this pure example is an assessment now of her deeply held and, as far as possible, stable preferences for trading off risk and return of long-term future outcomes. In other words, we need to assess her risk tolerance as an underlying psychological trait. In this example, nothing else matters. Fortunately there is a great deal of evidence that when examined using psychometric approaches, risk tolerance is both a measurable and stable trait.[8]

We cannot elicit the risk tolerance trait we need by observing investment choices, either real or through hypothetical questions put to investors; there are problems in extrapolating both long-term preferences from short-term choices[9] and aggregate preferences from choices over subcomponents.[10] Current choices may tell us quite accurately what risks people *do* take but not what risk trade-off we should optimise for when building their long-term portfolio. The latter is what is required from a risk tolerance measure.

CASE 2: INTRODUCING RISK CAPACITY

Now let's add one complication to the scenario and see how it differs from the situation above. Imagine that at the end of the 10 years, Christy will be *required* to purchase a single asset with the full amount of the portfolio as it stands at that date. This scenario, for example, was to some degree the case until recently in UK pensions, where many investors were compelled to purchase an annuity with their pension pot at retirement.[11] Does this change the risk that it is appropriate to take with her investments?

---

[8]For example, see Weber, Weber, and Nosić (2013); Davies and Brooks (2014); Sahm (2012); Egan, Davies, and Brooks (2011).

[9]For example, Benartzi and Thaler (1995); Samuelson (1963).

[10]For example, Thaler (1999); Rabin (2000).

[11]In reality the rules were more complex. For example, 25% of the value could be taken tax free as cash, and full compulsion for annuitisation was for those over 75.

The answer is yes, very significantly. This simple change sets up a future liability to be funded from the portfolio. It has no effect on Christy's risk tolerance, but it reduces her risk capacity. More importantly, it reduces her risk capacity in a dynamic way. Initially the effect on the appropriate risk level would be small. Ten years is a long way off, and initially almost all that would matter is her risk tolerance, as before. Christy would want to maximise her risk–return trade-off to get the best possible future outcome relative to her risk tolerance. As the date approaches, however, it would become increasingly important not to expose the capital value to the whims of the market. In particular, in the days leading up to the termination date, she would wish to reduce the risk level to zero to ensure she wasn't forced to sell at a low value just because the markets happened to have dropped that week or even that day.[12]

This, broadly, is the rational justification for target-date funds (TDFs), which—as originally conceived—gradually reduce risk to zero by a specific date (usually related to retirement).[13] However, there are two big problems with such funds. The lesser problem is that they tend to reduce risk linearly with time even though risk capacity doesn't reduce linearly but, rather, accelerates to zero as the final date approaches. The more important problem is that target-date funds are an ideal solution only for an ideal case: The moment the investor's situation becomes more complex than that of needing to fund a single asset purchase at a single termination date, such funds will deliver the wrong answer. The usual result is that most investors take substantially less risk than they should.

Note that in this scenario *nothing* about Christy's risk tolerance changes, and yet the appropriate solution changes dramatically. Indeed, over much of the life of the portfolio, the suitable answer will be driven much more by an assessment of risk capacity than by risk tolerance. Also, the risk capacity will change in a fairly systematic and measurable (indeed, to a large extent predictable) way with time, which means that the solution can be dynamically adjusted without any need for subsequent profiling.

The difference between these two situations can also be interpreted in terms of *time horizon* (another widely misunderstood notion in suitability). In the first case, Christy's true investment horizon is indefinite, despite leaving her assets with the investment manager for only 10 years; her investing is a going concern throughout, and beyond, the engagement.[14] In the second case, her investment horizon is 10 years initially and decreases steadily as that date approaches.

---

[12]Indeed, some of the biggest losers from the financial crisis were those exposed to risk assets in late 2008 and early 2009 and who were forced to annuitize this value during this time.

[13]In current practice, many TDFs never reach zero risk because the managers believe that even retirees need to take some risk—for example, to keep up with inflation.

[14]This point, however, raises a complication. The investment manager will want to deliver a good outcome at the point Christy will remove assets and thus will be incentivised to reduce risk as the 10-year point approaches, despite the fact that doing so is not at all in Christy's best interest.

CASE 3: MULTIPLE FIXED GOALS

Now let's add a further complication: Assume that every year along the way Christy will need to withdraw $50,000 from her portfolio. This sets up a chain of nine additional liabilities, which has a (negative) effect on her risk capacity. But, interestingly, it may also have implications for how we assess risk tolerance.

These liabilities may be seen as additional goals that Christy needs to achieve along the journey. Each is relatively small compared to the initial portfolio value of $1 million, but the higher the level of portfolio risk, the greater the chance that any given withdrawal will take a large chunk out of the portfolio if valuations are low at that point—making a good final value much harder to achieve.

As with the final asset purchase (a withdrawal of 100% of remaining portfolio value at the termination), each of these withdrawals will play a successively larger role in reducing risk capacity as they get closer to occurring. Each future liability will reduce risk capacity according to (a) how large it is and (b) how close it is in time.[15] If Christy has no anticipation of withdrawing from the portfolio to spend (as in Case 1), then her risk capacity is extremely high; if she needs to withdraw it all to spend tomorrow (as she would at the end of the 10 years in Case 2), then her risk capacity is near zero; for all intermediate cases, her risk capacity is reduced insofar as spending needs are larger and closer.[16]

Armed with this notion, it also becomes reasonably clear how we would adapt Case 1 to deal with any withdrawals Christy may plan to make *after* the end of the 10-year period. Such withdrawals will reduce her risk capacity, but because they are a long way off, their initial effect will be muted.

A big drop in portfolio value will mean that each subsequent $50,000 withdrawal becomes a larger percentage of the total portfolio, with the effect that these withdrawals will further reduce risk capacity.[17] In this scenario, the risk profile of the portfolio can still be adjusted automatically to provide the appropriate solution, without any recourse to further client profiling, as long as we can fully describe her financial circumstances and withdrawal needs. Everything we need to know about Christy in this scenario is defined by her circumstances (for risk capacity) and her risk tolerance.

---

[15]And, as we shall see in Case 4, also by how malleable this cash is with regard to timing, magnitude, or necessity.

[16]This concept would need to be expanded upon if we extended the example to allow Christy to have other assets and liabilities on her balance sheet and additional income and expenditure cash flows over time.

[17]An interesting corollary that follows from this notion is that a coherent model of risk capacity is very closely related to the ongoing debates around *sequencing risk* and should provide a useful tool in deriving appropriate solutions to this problem.

CASE 4: CONTINGENT GOALS

Now let's let Christy back in along the journey—but not fully. Let's still imagine she's immune to any emotional responses to her portfolio value over the journey— no doubt, no anxiety, no loss aversion, no panic—but she may decide to change her withdrawal amounts along the journey in response to how her portfolio is doing. If it is doing well, she may decide to increase the withdrawal beyond her planned $50,000; if it is doing badly, she may decrease the withdrawal or potentially forgo it altogether.

As well as providing a means to respond dynamically to changing circumstances, this example also enables Christy to place different levels of *priority* on these interim goals/ withdrawals. When she has no ability to modify or forgo interim withdrawals, we have in effect assumed that Christy sees each of these withdrawals as equally vital, which, because they occur earlier, implies that she gives them priority over achieving a good final value of the portfolio. When we allow her to modify or forgo these interim goals, we are instead implicitly saying that she may see these as less important than maintaining the final value of the portfolio.

This flexibility means she is able to treat her goals (or at least some of them) as *contingent*—similar to a charity tasked with trying to maintain the capital value of its endowment over time by varying its spending according to the realised portfolio value. This is a highly sensible way of varying the sequencing of withdrawals since, on balance, it means selling more when values are high and retaining more capital in the fund when values are low. This plan may be constrained by certain baseline spending requirements, but it is seldom the case that we have no ability to tighten our belts a little when circumstances require it.

As an alternative way of seeing this, imagine that Christy has vague plans to use a large amount of her portfolio—say, $500,000—to put toward a second home in 12 years' time. Should this goal feature strongly in reducing her risk capacity? Probably not. Because her goal is still quite fuzzy, she has plenty of opportunity to change it in response to future circumstances. If her financial position is less strong than she had hoped as the time approaches, she could choose not to buy a house, she could purchase a smaller, cheaper house, or she could put the plans on ice and delay the purchase until her portfolio recovers.

The ability to vary our spending according to circumstances provides psychological and financial *optionality*. The option value of this flexibility is extremely large. It increases risk capacity because we no longer have to assume that we must blindly meet every spending goal regardless of circumstances. In practical terms, we shouldn't count the full value of future liabilities that are low priority, flexible, or contingent; this ability to

adapt increases risk capacity. Many future goals will be contingent and low priority, and Christy should not dramatically shy away from taking risk now to protect capital for future plans that might not happen.

The ability to shrink or postpone goals if things don't turn out as we want means that we can afford to offer any future goal slightly less than full, unconditional protection. We thus make our finances more adaptable to changing circumstances, as well as to our own changing preferences and priorities. As things change, we can redeploy capital to those goals that are more important to us and shrink, delay, or forgo spending on goals that are no longer high priorities.

This is a major reason why we cannot truly assess risk capacity without embracing goal-based investing. It is also a major reason why most conventional approaches to goal-based investing,[18] which encourage investors to divide assets into "pots" to achieve individual spending goals, are sorely misguided. They encourage investors to be overly precise in their goals, and they reduce both the financial and psychological flexibility of being able to treat goals contingently and to transfer capital flexibly between goals as circumstances arise.[19]

Goals as a Component of Risk Profiling

But what about risk tolerance? This is the idea that the degree to which Christy's future goals affect her risk capacity depends on the importance she places on them and the degree to which she is prepared to postpone or shrink them, and it is a vital missing link in establishing her appropriate overall risk profile. It implies that we will never have a comprehensive description of the right investment risk level by just using a combination of risk tolerance and static risk capacity. Indeed, it is in the description of someone's overall *system* of goals, with all the valuable fuzziness and imprecision that this requires, that the notions of risk capacity and risk tolerance start to overlap and blur.

On the one hand, we have the concept of risk tolerance, which, as in Case 1, tells us how Christy views the trade-off of risk versus return for her indefinite-horizon investing—that is, investing that doesn't serve any specific withdrawal need and thus has no particular time horizon. On the other hand, we have the concept of risk capacity, which tells us how to decrease this risk level to protect future spending requirements depending on her weighted time horizon of future withdrawals[20] and how to change this risk depending on what else is on her balance sheet.

---

[18]Or attempts to apply liability-driven investment approaches to individual investors.

[19]See Brooks, Davies, and Smith (2015) for a behavioural perspective on goal-based investing.

[20]Or increase it if her future expected income flows outweigh her spending needs.

But between these we have the need to help Christy understand and articulate her system of goals, without imposing pressure for more precision than she feels is comfortable. This system of goals may be seen as a vital tool for establishing her risk capacity, accounting for the degree to which achieving each of her goals is essential to her. Without it, risk capacity won't provide an accurate description of the risk she can take.

But we can also see the description of her goals as a way of expressing more accurately her *tolerance* for risk: She should not be prepared to tolerate risks that place her coveted aspirations in danger or that place her cherished assets, her already achieved *crystallised goals*, at risk. Her risk tolerance could be reasonably viewed as an expression of the degree to which she is prepared to place a higher-priority goal at risk in order to have the chance of achieving both this more important goal and some other, less important goal as well. These may not be goals she can yet articulate or express, but ultimately, taking investment risk is a way to expand the set of financial goals one can attain while exposing oneself to the risk of shrinking this set. So risk tolerance, too, is about goals in the end.

The logical conclusion from this line of reasoning is that if we genuinely want to understand the right amount of risk for an investor to take, it is not sufficient just to measure risk tolerance. It is also not sufficient to supplement this with a narrow model of risk capacity. We fundamentally also need to help people understand, articulate, and dynamically adapt the totality of their future goals, plans, and aspirations over their journey.

Effective goal-based investing is the truly new vista in risk profiling.

CASE 5: BEHAVIOURAL ASPECTS OF RISK ATTITUDE

The last step we will take is to make Christy fully human. She now notices (sometimes) what is happening to her portfolio value, she concerns herself (sporadically) with the details of what her portfolio is invested in, and she has (frequent) behavioural and emotional responses to what happens over the investment journey. All of these cause her experienced risk attitudes to fluctuate at various points in time. Sometimes she acts on these fluctuations.

These *behavioural risk attitudes* may be considerably at odds with her underlying *risk tolerance* but nonetheless reflect how she feels about the risks she is taking at any given point in time. Christy, like all of us, is subject to anxiety, doubt, exuberance, myopia, and framing effects. It is impossible to deny that all these behavioural responses are a large part of her risk attitudes. It is also impossible to deny that these influence her decisions.

Understanding these ephemeral risk attitudes is thus an important part of any comprehensive risk-profiling exercise and should be used to guide investment advice. However, these behavioural aspects, revealed through observed behaviour, should not be used as the foundation for an assessment of risk tolerance and portfolio optimisation. Why would you "optimise" a risk–return trade-off for someone on the basis of all the error-prone, bias-ridden, context-dependent, pro-cyclical, unstable, myopic, and inconsistent behaviour they exhibit? Doing so is committing the *naturalistic fallacy*: You can't derive "ought" from "is." For example, we cannot infer from the fact that investors *do* exhibit loss aversion that these attitudes are what they *ought* to use to optimise their portfolio.

Risk tolerance is about Christy's willingness to accept risky *outcomes*, whereas these other behavioural aspects are more about her shorter-term emotional resilience along the journey. This distinction has huge ramifications for determining the correct investment solution and the degree to which risk attitudes should be built into portfolio solutions on the one hand or into the client relationship on the other. Risk tolerance, combined with her risk capacity, tells us what level of risk Christy *should* strive for in her portfolio. Her other behavioural risk attitudes instead affect her emotional comfort with this solution over the journey.

This doesn't mean we should ignore these risk attitudes, but we shouldn't pander to them either. Managing anxiety and emotional comfort along the journey are at least as important in attaining a good outcome as determining the "optimal" long-term solution. For example, using better client communication and long-term framing of investment performance reports may mitigate the emotional effects of loss aversion much more effectively and cheaply than changing the investment strategy to avoid instances of loss along the journey. No one will be completely comfortable with the compromises that have to be made in any investment program, so we should instead seek to make Christy *sufficiently* comfortable with a *sufficiently* efficient solution.[21]

We may draw a useful analogy between Christy's behavioural risk attitudes and her risk capacity. Just as risk capacity is affected by future financial liabilities that reduce her weighted time horizon of withdrawals, so her behavioural tendencies create the potential for future emotional liabilities (withdrawals) that reduce her emotional time horizon. Our human tendency to myopia and behavioural responses to our immediate environment means that in most cases we end up making decisions based on an emotional time horizon that is way shorter than our financial time horizon.

---

[21]See Davies (2014) for an extended discussion on how these behavioural tendencies can be measured using psychometric approaches and how these can be used to practically improve investors' emotional comfort.

As Keynes is reputed to have said (although there is no evidence that he actually did), "Markets can remain irrational for longer than you can remain solvent."[22] This is essentially an injunction to consider risk capacity when investing. The need to withdraw assets from your portfolio to meet a future liability (i.e., the need for financial liquidity) increases the risk of market values being "irrationally" low when this time comes; that is, it reduces risk capacity. However, many, if not most, of those who sell their risky assets at the bottom of a financial crisis do not do so because they have run out of financial liquidity. Most do not, at that single point in time, need their entire portfolios to be liquidated to fund spending. In other words, they are not insolvent; they still have risk capacity. Instead, they run out of *emotional liquidity*. They are afraid they will lose the rest of their money. Immediate behavioural risk aversion spikes, and emotional time horizons decrease to zero.

The crucial difference is that the shorter time horizon that comes with reduced risk capacity is one we *should* reflect in the optimal portfolio. Take more risk than you have capacity for and you risk financial insolvency. The shorter time horizon that stems from behavioural risk attitudes, in contrast, is one we should fight against, not embrace. Changing "optimal" portfolio risk to reflect emotional time horizons and behavioural risk attitudes takes investors away from the appropriate solution, not toward it. Of course we can't ignore the risks of emotional insolvency; when we are stressed, we crave emotional comfort at any cost and hence make costly investment decisions to acquire it. But pandering to these emotional responses is a very expensive way of sleeping better at night.

The correct solution for the *Homo economicus* version of Christy that is portrayed in Cases 1–4 is the conventional goal of finance theory: to maximise her *risk-adjusted returns*, where the adjustment for risk derives from the appropriate risk profile that results from the combination of her risk tolerance and risk capacity. The right solution for the more human Christy of Case 5 has the *same risk profile*—she has the same long-term needs and preferences—but the goal should instead be to maximise *anxiety-adjusted returns*.

The information we gather in profiling her behavioural risk attitudes is vital in understanding what makes her anxious and uncomfortable along the journey so that we may help her overcome these adverse influences on her actual risk taking. With this insight, we can help her buy the emotional comfort she requires as efficiently and cheaply as possible—that is, with as little deviation from the solution that fits her long-term risk profile as possible.

This is not to say that behavioural risk attitudes should never lead us to change the suitable risk levels, asset allocation, or individual investments in her portfolio. These may all be valid ways of reducing anxiety. For example, her ambiguity aversion might be

---

[22]This quote is frequently attributed to Keynes but was first recorded in the early 1990s.

assuaged by deliberately introducing some home bias, or small concentrations to familiar assets among her holdings. Doing so might be suboptimal relative to the theoretically "perfect" portfolio, but this small loss of efficiency might be well spent if it makes Christy more comfortable with investing over the long term. Similarly, dollar-cost averaging is technically inefficient and thus theoretically "wrong,"[23] but in many cases, it could be a very cheap way of giving someone the minimal emotional comfort they need to get invested at all. Maximising anxiety-adjusted returns can even be brought into a quantitative portfolio optimisation process that targets specific measures reflecting those aspects of portfolio paths that are most likely to induce anxiety.[24]

But the crucial point is that satisfying behavioural risk concerns are never the *goal* of investing and, as such, should never be used as inputs to construct the "optimal" portfolio: They are what we need to *overcome* in order to get closer to the optimal portfolio. And sometimes that means moving slightly away from the "best" solution if by doing so we purchase a big reduction in anxiety.

Do not let the best be the enemy of the good.

MAPPING THE LANDSCAPE

Bringing all the cases together, we can now discern a mapping between four reasonably distinct strands of an investor's overall situation and different aspects of his resulting profile, as shown in **Exhibit 1**. These four strands fall roughly on a continuum from purely financial to purely behavioural:

1. Financial circumstances—the investor's current balance sheet, cash flow expectations, and insurance coverage

2. Goals and aspirations—the investor's future spending goals

3. Personality traits—long-term, stable proclivities toward certain behaviours, such as risk tolerance or innate investment composure

4. Emotions and behaviour—short-term behavioural responses to context and environment

This "continuum" maps nicely to the traditional categories of *ability to take risk* (risk capacity) and *willingness to take risk*. The former largely reflects financial circumstances,

---

[23]On the grounds that leaving wealth uninvested when expected risk-adjusted returns are positive is suboptimal; having determined an optimal asset allocation, an alternative allocation that dilutes this with cash cannot also be optimal.
[24]See Davies and Lim (2014).

whereas the latter is more about personality. However, as discussed, future cash flows and goals can play a pivotal part in defining both ability and willingness.

We may characterise risk tolerance as the investor's long-term, stable personality preference for risk versus return. Alternatively, we can think of risk tolerance as the investor's willingness to risk attaining fewer future goals for the chance of attaining more future goals. This framing in terms of goals is helpful but not necessary: Some investors may, more simply, prefer to think of risk tolerance as the willingness to risk a certain amount of money for the possibility of earning even more money.

However, willingness to take risk also includes numerous behavioural risk attitudes: short-term behavioural responses (among others, reference dependence, framing, loss aversion, ambiguity aversion, risk perception, myopia, hyperbolic discounting, and probability distortion) that reflect both the immediate context and aspects of the investor's personality (such as her innate degree of investment *composure*).

EXHIBIT 1. INVESTMENT PLANNING USING RISK CAPACITY AND RISK TOLERANCE

Financial Circumstances	Goals and Aspirations	Personality Traits	Emotions and Behaviour
Ability to take risk	Willingness to take risk		
Risk capacity	Risk tolerance	Behavioural risk attitudes	
Investor's *risk profile* – seek to satisfy these requirements		Seek to mitigate, not appease	

Exhibit 1 sums up these ideas. The final row shows that risk capacity and risk tolerance together form essential components of the appropriate (or suitable) risk level for an investor's overall investable assets. Behavioural risk attitudes, however, while also essential to understand, reflect not the risk an investor *should* take but, rather, the emotional currents that constantly pull investors off the appropriate path along the investment journey. They need to be mitigated, neutralised, or overcome—and sometimes, to as small a degree as possible, appeased. But they should never be mistaken for the right answer.

Much of the confusion surrounding current risk-profiling practices stems from a lack of clarity about the boundaries between these concepts—in particular, the following:

1. Undue focus on risk tolerance at the expense of understanding risk capacity (measuring what can be measured most easily), making risk tolerance do more of the work than it should in determining suitability

2. Confusion between risk tolerance and behavioural risk attitudes, with a tendency to want to treat observations of the latter as a proxy for the former

3. A lack of any clear model of risk capacity that truly accounts for the full range of influencing factors—specifically, a failure to understand the crucial role that investor goals play in an adequate description of risk capacity or the fundamental importance of imprecision, fuzziness, and adaptability of these goals

4. Failing to realise that the true suitability of an investment solution can be demonstrated only in the context of the investor's overall situation—neither risk tolerance nor risk capacity can be intelligibly measured for isolated components of an investor's wealth

NEW VISTAS

So given this landscape, how can we better set about the practical process of effective risk profiling?

Comprehensive descriptions of risk capacity models and goal-based investing would each require an article of their own (at least). However, it is fair to say that each of these requires vastly greater attention than they currently receive. Here I shall focus on the narrower question of measuring attitudes.

MEASURING RISK TOLERANCE

A corollary to the lack of focus on risk capacity is that little would be lost by sticking to simple psychometric approaches to measuring risk tolerance. There is an unnecessary tendency to want to overcomplicate this easiest-to-measure aspect of an investor's risk profile. Efforts to pour more resources into ever more sophisticated ways of eliciting risk tolerance are to a large degree unnecessary and in many cases misguided. There is a danger of spurious precision in risk tolerance measurement, and in striving for greater precision, there is a concern that we measure revealed preferences and observed behaviours that are more properly reflective of behavioural risk attitudes than of true risk tolerance.[25]

Measured correctly, decent psychometric measures of risk tolerance are stable and differentiate effectively and reliably between individual investors. The gap lies not in how risk tolerance is measured but in how it is used. A simple, well-designed measure, used well, is hard to beat.

---

[25]A discussion of existing methodologies for eliciting risk attitudes may be found in Egan, Davies, and Brooks (2011).

That said, there are enticing opportunities for measuring risk tolerance in ways that offer a better client experience and greater accuracy. One is making existing psychometric questionnaires more dynamic, making the questioning process more responsive to existing client knowledge and initial responses, and gleaning more information out of fewer questions.

This *might* be supplemented with asking clients to respond to explicit hypothetical choices (for example, presenting them with choices between investment options or combinations of risk and return) and visual stimuli, using pictures rather than questions. However, such approaches often give very momentary readings of how an investor responds to specific questions and may not offer generalizable insights, instead providing only overly precise estimates that hold true only in these hypothetical, and often highly artificial, situations. It will be crucial to validate the idea that such tools correlate strongly to existing psychometric risk tolerance measures.

A further area of promise is to use big data, observed behavioural patterns, and machine learning to home in on an investor's risk tolerance. Such data have been very effective in assessing numerous personality traits,[26] and since risk tolerance is best understood as a personality trait, it could well lend itself to these techniques. However, there are valid concerns about data protection with many of these applications, and in deploying such techniques, it is crucial that we are able to clearly discern between risk tolerance and behavioural risk attitudes.

On balance, anything more complicated than a simple well-designed and validated psychometric questionnaire is likely to be both unnecessary and less robust.

MEASURING BEHAVIOURAL RISK ATTITUDES

Unearthing unstable behavioural risk attitudes is where observing real behaviour has the greatest potential. Behavioural responses are complex, individual, and unpredictable. As long as methods to measure them are used thoughtfully and intelligently, observing patterns of changing attitudes to risk may well help in guiding investors to better behaviour along the journey. These profiling techniques may include both psychometric instruments and observing historical behaviour. Each provides valuable insight on what makes investors anxious and uncomfortable. By examining how investors really respond to a wide range of actual situations, we can start to build a rich profile of their behavioural tendencies and identify spots where these are likely to be costly.

---

[26]For example, the myPersonality project of the Cambridge University Psychometrics Centre, in which correlations between individuals' online Facebook profiles and their personality traits are studied.

With knowledge of the investor's full balance sheet and armed with systems describing the cash flows and (imprecise) goals required for risk capacity, another option becomes possible. Rather than just looking at history, which exposes investors to only a limited range of all the possible events they might have to bear, we could instead expose investors to simulated versions of their future financial lives and use responses to these as profiling inputs. Even more valuably, such risk simulations have been shown to improve long-term investment decisions[27] and thus can be used to educate and prepare investors for the future.

This points to where profiling is potentially most valuable with regard to behavioural risk attitudes: Profiling should provide an opportunity for investors to learn about their attitudes, emotions, and biases, and in doing so, it should help them prepare for the anxiety of the journey.[28] The more this can be done using investors' own past behaviour and their specific financial situation and goals, the more effectively it will help them to take and stick with the right amount of risk.

DYNAMIC PROFILING

Possibly the most exciting new vista in risk profiling is not about the specifics of what is being measured; rather, it is in when and how the data are collected and used. At present, suitability, client fact-finding, and profiling approaches can be off-putting. The conventional presumption is that a comprehensive profile is needed from the outset, so clients are often exposed to a great deal of questioning up front. This situation results in a poor client experience and poor subsequent commitment, as well as in advisers cutting corners and "box-ticking" in order to hasten the process. Little is gained in terms of client education or emotional preparation for investing. And the solution arrived at is then often over-fitted to the client's circumstances, opinions, and emotional state at that one specific point in time. This solution is then regarded as the right ongoing solution, sometimes for many years, despite constant shifts in the investor's portfolio, circumstances, preferences, and goals.

Instead, we need to use advances in technology, data analytics, and behavioural design (in particular, techniques from *gamification* to help investors engage comfortably with complex systems[29]) to blur the current distinctions between *profiling*, *suitability*, and *client engagement*. The initial process should be as minimal, simple, and easy for the investor as possible. The more effort it takes to start investing, the more investors

---

[27]Bradbury, Hens, and Zeisberger (2016).

[28]"Beyond Risk Profiling" (2015).

[29]Note that by "gamification," I do *not* mean turning the process into a game or trivializing it. Instead, gamification is, or should be, the use of tools and techniques (often borrowed from, or used in, the gaming industry) to lead people rapidly and comfortably to sophisticated patterns of behaviour that enable them to engage effectively with complex environments or problems.

(particularly first time investors) are put off. However, after starting, advisers should be constantly profiling.

Every engagement point is an opportunity to enrich the profile—to ask the few next questions that offer the most value in deepening customer understanding. Every change in the profile offers an opportunity to sharpen the overall advisory solution: either directly through changing the portfolio (although here the range of options to respond to small changes in the profile is relatively constrained) or indirectly by changing how we engage and communicate with the client, how we manage clients' emotional states over the journey, how we plan and program future actions and responses, how we report performance, and how we involve the client in ongoing decisions. And every change to the overall solution is an opportunity to identify new valuable engagement opportunities. This approach conceives of risk profiling as part of a never-ending dialogue between investor and adviser, constantly updating the profile to gather more information and adapt to changes in circumstances and preferences.

The risk tolerance measure at the heart of the profile remains the steady core around which the profile is built and will in all likelihood still be most effectively arrived at through a simple psychometric questionnaire at the start of the investment relationship (and which is refreshed occasionally). But once the basics are in place and clients have invested, their risk capacity, goals, and behavioural risk attitudes can become an entirely dynamic process, both adapting and deepening through the client relationship. Of course this is much how human advisers' subjective understanding of their clients' needs deepens over time. But with advanced data analytics, gamification, and machine learning, we are coming to the point where we can blur the distinction between formal profiling and an adviser's client understanding, ultimately creating a profiling process as a digital *decision prosthetic* to supplement and enhance the essential human side of investing.

In conclusion, if we want to move forward in risk profiling as an industry, we should resist the temptation to focus on ever more complex batteries of questions to elicit and parameterise complex approaches to risk tolerance. Simple is best. We should instead focus on building and improving dynamic models of risk capacity and tools to truly help investors understand and articulate their own goals without requiring that they express these goals with any greater precision than that with which they are truly held. We should concentrate much more on how we *use* this knowledge and include simple risk tolerance measures in coherent suitability frameworks that reflect an understanding of what truly matters to clients. And, finally, we should stop treating profiling as a single point-in-time activity, divorced from the ongoing client relationship, but instead ensure that suitability adjusts dynamically to meet constantly changing needs.

The ultimate new vista in risk profiling is that it becomes impossible to separate from goal-based suitability and effective client engagement.

REFERENCES

Benartzi, S., and R. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics*, vol. 110, no. 1: 73–92.

"Beyond Risk Profiling: Achieving Better Investment Outcomes for Consumers and Industry." 2015. White paper, True Potential Centre for the Public Understanding of Finance (October).

Bradbury, M., T. Hens, and S. Zeisberger. 2016. "Do Risk Simulations Lead to Persistently Better Investment Decisions?" Working paper (12 July).

Brooks, P., G.B. Davies, and R.E.D. Smith. 2015. "A Behavioral Perspective on Goal-Based Investing." *Investments & Wealth Monitor* (November/December).

Davies, G.B. 2014. "Overcoming the Cost of Being Human." White paper, Barclays.

Davies, G.B., and P. Brooks. 2014. "Risk Tolerance: Essential, Behavioural, and Misunderstood." *Journal of Risk Management in Financial Institutions*, vol. 7, no. 2 (Spring): 110–113.

Davies, G.B., and A. de Servigny. 2012. *Behavioral Investment Management: An Efficient Alternative to Modern Portfolio Theory*. McGraw-Hill Professional.

Davies, G.B., and A. Lim. 2014. "Managing Anxiety to Improve Financial Performance: Don't Let the Best Be the Enemy of the Achievable." In *Quantitative Approaches to High Net Worth Investment*. Edited by A. Rudd and S. Satchell. London: Risk Books.

Egan, D., G.B. Davies, and P. Brooks. 2011. "Comparisons of Risk Attitudes across Individuals." In *Wiley Encyclopedia of Operations Research and Management Science*. Wiley.

Odean, T. 1999. "Do Investors Trade Too Much?" *American Economic Review*, vol. 89, no. 5 (December): 1279–1298. doi:10.1257/aer.89.5.1279

Payne, J.W., J.R. Bettman, and D.A. Schkade. 1999. "Measuring Constructed Preferences: Towards a Building Code." *Journal of Risk and Uncertainty*, vol. 19, no. 1–3: 243–270. doi:10.1023/A:1007843931054

Rabin, M. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica*, vol. 68, no. 5: 1281–1292.

Sahm, C.R. 2012. "How Much Does Risk Tolerance Change?" *Quarterly Journal of Finance*, vol. 2, no. 4 (December).

Samuelson, P.A. 1963. "Risk and Uncertainty: A Fallacy of Large Numbers." *Scientia*, vol. 98 (April/May): 108–113.

Slovic, P. 1995. "The Construction of Preference." *American Psychologist*, vol. 50, no. 5 (May): 364–371.

Thaler, R.H. 1999. "Mental Accounting Matters." *Journal of Behavioral Decision Making*, vol. 12: 183–206.

Weber, M., E.U. Weber, and A. Nosić. 2013. "Who Takes Risks When and Why: Determinants of Changes in Investor Risk Taking." *Review of Finance*, vol. 17 (July): 847–883.