

5

Don't Let the Best Be the Enemy of the Achievable: Managing Anxiety to Improve Financial Performance

Greg B. Davies, Antonia Lim

Barclays

For decades theoretical finance has sought to provide a normative solution, the "right answer", to the question of how to construct an investment portfolio. Unfortunately, the answer is the optimal choice only for a particular species of investor that does not exist in reality: *Homo economicus*, the "economically rational" investor. As the solution for real investors, modern portfolio theory (MPT) fails to account for much of what makes us human: our broad range of objectives and preferences, our psychology and our emotional responses to the gains and losses of our investments over time, ie, how we feel along the investment journey. MPT gives investors sophisticated machinery, but no instruction manual; in the hands of real people who cannot emulate *H. economicus* this can be more harmful than beneficial.

For MPT to work in the real world, we need to adapt these normative solutions to ensure they are the optimal solutions for real investors, for whom the pursuit of some long-term "right" answer may lead to emotional discomfort, and anxiety. We shall use the term "anxiety" to refer broadly to any short-term behavioural responses to investors' immediate circumstances, whether good or bad, that might make them uncomfortable enacting, or sticking with, the long-term normative solution. The truly optimal portfolio for real investors will require predicting, and insuring against, the kinds of anxieties that tempt investors off the path to long-term returns. There is nothing rational about pursuing a theoretically optimal solution knowing you will fail to enact it effectively. The rational solution embraces our own behavioural fallibility, and plans for it.

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

While these concerns are by no means limited to high net worth individuals, the relative importance of continual and consistent sound investment decision-making is much higher for wealthy individuals than others. This is because in general the ratio of current wealth to income and expenditure flows is much higher for the wealthy. For those without a large stock of wealth, sound investment decisions are much less important for overall financial efficiency than are decisions around budgeting, insurance, financial planning and savings. For high net worth individuals, with occasional exceptions, these questions around managing income and expenditure flows matter less to overall holistic wealth management than how well they invest the existing stock (ie, total investible amount) of wealth.

In this chapter we examine how investors should rationally deviate from the normative model in order to most cheaply reduce anxiety, and thereby attain higher realised returns from investing. In the next section we briefly outline the conventional normative solution. We then consider various reasons why this may not provide the optimal solution for real investors to pursue, concluding that preferences for short-term emotional comfort require investors to deviate from the normative model if they are to have any practical hope of achieving their financial objectives. Reducing anxiety mitigates costly behavioural responses to the environment along the investment journey. We discuss what we mean by anxiety, and how we might be able to precisely define and measure it. The fifth section incorporates the most promising measure (a version of drawdown) into an empirical exploration of methods of reducing anxiety when constructing portfolios, seeking the greatest anxiety reduction for the smallest cost in terms of pure financial efficiency. The final section concludes.

THE TRADITIONAL NORMATIVE CRITERION

The traditional objective of MPT is the maximisation of risk-adjusted returns. It provides a model to determine the portfolio that will attain the highest returns relative to the risks taken, given investors' financial preferences. Theoretically, this is accomplished by specifying the utility function that ascribes a utility value to each possible financial outcome, although in MPT this may be distilled into the investor's optimal trade-off between risk and return of the anticipated outcome

distributions for possible portfolios at the end of the investment time horizon. Both are governed by the investor's degree of risk tolerance.

For the normative solution to be genuinely optimal, it is necessary that our future beliefs (estimated distributions of future outcomes) reflect the reality of future possibilities as closely as possible, and that the utility function (or the investor's risk–return trade-off function) is an accurate expression of the investor's preferences. Estimating the future distributions of returns on which to maximise utility is the subject of reams of research (see, for example, McNeil *et al* 2005). Ensuring the utility function embodies accurate investor preferences has been investigated far less. Indeed, the quadratic utility function assumed as a basis for MPT is decidedly unreasonable as an expression of normative preferences for risk and returns (Davies and de Servigny 2012, Chapter 2). This issue is usually assumed away in practice by not specifying risk–return preferences at all, but instead determining a set of portfolios on the risk–return efficient frontier that contains the optimal portfolio(s) for multiple degrees of risk aversion (although picking one of these portfolios still requires some assumption about risk tolerance to be made).¹

Once we grant that both beliefs and preferences have been accurately expressed, the resulting portfolio is the normative choice for that investor. *Homo economicus* would make this choice, which maximises the expected utility of the returns distributions at the appropriate horizon, given the investor's degree of risk tolerance, T , or equivalently by maximising the utility of their trade-off between risk (R) and expected return (\bar{r})

$$\max \mathbb{E}[u(r), T] = \max f(\bar{r}, R, T)$$

IS NORMATIVE RATIONAL?

Considering non-financial objectives

Is this normative portfolio the rational choice in reality? The answer is a clear "no", and not because the normative model is wrong, but because no real investor is *Homo economicus*. MPT assumes that investors are very single minded in their objectives, caring only about financial returns. In reality, however, they have many additional non-financial objectives when investing. The optimal risk–return portfolio may not be the portfolio that best serves the totality of the investor's preferences.

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

In one regard, this difficulty is rather trivial. Of course individuals have other objectives in life, and of course the level of financial returns from their portfolio is not the sole thing they care about. Indeed, for most investors the objective of greater wealth is itself largely a means to other ends. One approach taken extensively in the literature is to use consumption, rather than wealth levels or returns, as the carrier of utility (Samuelson 1969; Merton 1971). However, this adds greatly to the difficulties of specifying the problem adequately to arrive at practical investment solutions. Preferences for consumption of various kinds are constantly in flux, and the paradigm of "constructed preferences" (Slovic 1991) leads us to doubt if they coherently exist at all.

For many, wealth may not be exclusively a means to an end: the level of wealth may itself yield utility of a kind. Alternatively, utility may be derived from the process of investing itself, with benefits that are not measured by financial return of the portfolio. For example, realisation utility (Barberis and Xiong 2011; Frydman *et al* 2014; Ingersoll and Jin 2013) comes from bursts of utility at the moment of sale rather than from the level of wealth or subsequent consumption. Furthermore, the act of trading may provide excitement, enjoyment or entertainment (Dorn and Sengmueller 2009; Kumar 2009). Certain investments, such as hedge funds or art funds, may confer status in addition to their financial returns, and, increasingly, many investments (ethical, social impact, etc) are promoted in terms of the social good they deliver in addition to the purely financial benefits to the investors (see, for example, Statman 2010). This would imply that the optimal solution would be to optimise a utility function with multiple arguments, reflecting utility from different sources.

But even if investors glean all manner of side benefits from investing, the advantage of the normative solution is that we can abstract from these complications and simply ask: if our objective was to maximise the financial returns taking into account risk aversion, what would the best solution be? The investor could then reasonably decide to sacrifice some financial efficiency for the emotional benefits of excitement, status or social impact as they see fit, but this does not imply that we should incorporate these into the normative solution itself.

Given all the complexity and instability of non-financial preferences, for practical investment purposes it is better to treat wealth

explicitly as a means to other (unspecified and varied) ends, and focus on helping the investor to construct coherent preferences for future returns. The normative model has value as a guide to help investors maximise their ability to fund future preferences, whether consumption based or not, without attempting to prejudge what these might be. It provides investors with a clear answer to the more precise and tractable question of choosing the best portfolio from a purely investment perspective: the financially rational solution to pursue. Investors may then choose to incorporate more subjective preferences as they see fit by choosing to deviate from this financially optimal solution by sacrificing financial efficiency to satisfy non-financial preferences.

Preferences for a comfortable journey

However, there is a distinct class of common non-financial preferences that have deep implications for rational portfolio choice. As investors, we have behavioural responses to outcomes along the journey. We respond emotionally to interim returns; we experience disutility from anxiety or discomfort in response to volatility along the journey; and we have preferences for feeling comfortable at each point in time with the risks to which we perceive our wealth to be exposed. Many of these responses arise from short-term impressions of the immediate context and environment, or from behavioural responses to the individual history of the past experiences of the investor, both of which are irrelevant to the optimal portfolio decision for the future.²

For example, because of empirically observed reference dependence in non-expected utility preferences, such as those modelled in prospect theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992), different paths of recent experiences will induce immediate reference-dependent responses in investors that will affect how comfortable they feel about the portfolio they hold. Recent experiences or the current context may also mean that the investor distorts perceived risks away from previous estimates, and finds it difficult to stick with previous assumptions. This may be useful where the evidence provides reasons to amend the distribution of outcomes at the decision horizon (interim points at which the portfolio is reassessed against its objectives and future investment environment), but frequently these changes in risk perception are highly unstable (see,

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

for example, Weber *et al* 2012) and reflect distorted perceptions and biases driven by current media concerns and focus, rather than any genuine change in the original underlying assumptions. This can make it difficult to stick with optimal investment plans due to short-term, and short-lived, perceptual biases, framing effects and behaviourally driven emotional discomfort.

In short, in addition to our financial preferences, we have short-term preferences for a comfortable ride along the investment journey and to assuage behavioural distortions caused by the immediate context and environment and past experiences. This may at first appear to be no different from the non-financial preferences identified above: like the pursuit of excitement, status or social impact, if investors have strong preferences for short-term emotional comfort, should they not just choose to deviate from this normative solution as they see fit? In other words, like all other non-financial preferences, why should we not just exclude emotional preferences for comfort during investment from what is considered the rational solution?

The answer is that to a large degree investors cannot avoid their need to reduce anxiety along the journey. Behavioural responses to the journey lead to biases of perception, to stress and to anxiety along the journey, which will naturally lead investors to deviate from financially optimal decisions along the journey, whether they like it or not. This leads to behaviour that is financially costly, even if the investor’s intent is to maximise only the expected utility of financial returns.

The more we feel emotionally uncomfortable, stressed, anxious (or, conversely, elated and exuberant), the more our behaviour will be driven to satisfy these short-term urges. The framework of two systems of reasoning (Sloman 2002; Kahneman 2011) suggests that our behaviour will frequently be driven by choices that “feel” right and are influenced by framing, context and the intuitive responses of our impulsive, emotional “system 1”. Our deliberative “system 2” reasoning may help us to overcome “hot cognition” along the journey, but even deliberation will not necessarily arrive at a normatively correct choice (we are innately poor statisticians, even when thinking actively), and we can neither turn these impulses off nor avoid to some degree pursuing emotional comfort at the cost of deviations from financially normative decisions.

The cost of behavioural preferences

This chapter will not extensively survey the abundant evidence for such deviations,³ but we note that these costs broadly divide into two categories: anticipation of discomfort prior to investing; and desire to increase comfort having invested. First, the anticipation of discomfort, the fear of regret that the investment will perform badly (Loomes and Sugden 1982), excessive focus on the possibility of future losses, or of unlikely but psychologically salient negative outcomes, might make the investor reluctant to undertake the financially optimal investments in the first place. This may result in systematic underinvestment, with too much capital held in cash. Or it may lead the investor to enact only those components of the optimal portfolio that feel sufficiently comfortable. Usually this means forgoing, or perennially postponing, the riskier components of the normative portfolio, or undertaking only those investments that come with a comfortable back story (good media coverage, familiar assets, home country assets, funds with good recent track records, assets that have connections to friends and family or simply those with the best marketing). This reluctance is costly, because it results in either a suboptimal risk–return trade-off, or over-concentration in convincing narratives, which are seldom the same as good investments (for example, near-ubiquitous home bias in portfolios (French and Poterba 1991)).

This is likely to be exacerbated by our tendency towards “narrow framing” (Thaler 1999): approaching decisions in isolation from the bigger picture. We exhibit this both over time, by focusing myopically on the short term rather than aligning our decision-making to the longer time frame of our objectives (for example, the “myopic loss aversion” of Benartzi and Thaler (1995)), and at a single point in time by considering each investment in isolation rather than in the context of the overall portfolio. The danger is that these isolated and short-term decisions may not aggregate to a portfolio that is in any way optimal for the investors’ long-term financial preferences regarding their total wealth. Rabin (2000) shows that the risk aversion we observe over small amounts results in absurd levels of risk aversion when applied to total wealth, linking this to behavioural effects of reference dependence and loss aversion; the implication is that, because of narrow framing, individuals will generally make individual investment choices that aggregate to a level of risk for

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

their overall portfolio that is considerably less risky than the optimal risk–return trade-off.⁴

Second, once they have invested, investors will act in ways that increase their comfort over the journey at the expense of deviations from their financially optimal strategy. They may find themselves doubting their initial long-run assumptions after reading emotionally charged media articles, and be tempted to reduce risk or shift away from certain asset classes in pursuit of perceived short-term safety. They may cling too tenaciously to losing assets, and too easily take profits on winning assets, showing the “disposition effect” (Odean 1998; Genesove and Mayer 2001). They may also be driven by an action bias to seek comfort in making constant changes to the portfolio, which – even if not driven by the pro-cyclical risk taking that leads to a “behaviour gap”,⁵ or costly switching to chase funds with recent good performance (Frazzini and Lamont 2008) – still results in reduced returns from transaction costs (Barber and Odean 2000).

The existence of reference-dependent loss aversion along the journey will trigger short-term, frame-dependent distortions. Two investors facing the same future probability distribution, and with the same normative utility function, will nonetheless feel differently about this distribution depending on the placement of their current personal reference point, which will differ according to the path by which they arrived at this point. For example, if one of the investors has recently lost money and requires returns greater than 5% before arriving back at the level they perceive to be neutral, all returns below 5% will be psychologically coded as losses. The other investor may instead have recently gained a windfall increase in wealth, and will psychologically respond to actual losses as though they are merely diminished gains. This version of the house money effect (Thaler and Johnson 1990) will lead to a different emotional evaluation of the same probability distribution from the first investor, who instead faces enhanced loss aversion.

We cannot avoid these costs by ignoring them, although this is the approach taken by classical finance: ignore emotional responses, or rather fail to recognise their existence, and pursue the theoretical optimum provided by MPT. It is not rational to pursue a solution, no matter how normative, that fails to incorporate knowledge of our own fallibility. By assuming the financially normative solution is also

the rational solution to pursue, we set ourselves up for expensive failure.

Purchasing emotional comfort, efficiently

There is no question over whether we can suppress all need for short-term emotional comfort: we cannot. Instead, we should recognise that a certain amount of emotional comfort is essential if we are to get anywhere near the theoretical solution. The right question is how much comfort is sufficient to enable investors maintain a strategy as close as possible to normative. How do we acquire maximum emotional comfort (or, equivalently, anxiety reduction) at the least cost to the financially optimal solution? By deviating from the normative solution in ways that buy the greatest amount of satisfaction of our short-term behavioural preferences, we can reduce the likelihood of unplanned, knee-jerk and expensive behavioural reactions along the journey. Instead, we purchase insurance against our own likely emotional responses in a planned, cheap and efficient manner.

Thus, if the normative solution requires maximising the expected utility of a point-in-time distribution of future outcomes, the behavioural solution involves recognising that this solution is unattainable in reality: we are human, emotional and fallible. So even where we wish to strive for the best solution to satisfy purely financial preferences, we still need to assuage our short-term behavioural need for emotional comfort over the journey. To formalise this, we can calculate the normative expected utility of possible portfolios, but then note that the possible journeys taken by some portfolios will be more likely to induce anxiety than others. We can thus characterise each portfolio by its expected utility and a measure of anxiety.

Expected utility is a function of the expected return (minus the compensation required for the risk) of the distribution of portfolio returns at the point in time at which the investor next intends to (or is able to) make active changes to the portfolio.⁶ In the formulation of Davies and de Servigny (2012) optimising expected utility is equivalent to maximising desirability. Desirability (D) is defined as

$$D = (\bar{r} - r_f) - \sigma_B^2/T$$

where the first term in the parentheses is the expected excess return of the portfolio over the risk-free rate, the second term is the minimum compensation that an investor with a particular degree of risk

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

tolerance, T , will require for taking that risk.⁷

$$\sigma_B^2 = \frac{T^2}{2} \ln \mathbb{E} \left[\exp \left(\frac{2(\mathbb{E}[r] - r)}{T} \right) \right]$$

is a behavioural generalisation of variance, which accounts for the investor’s normative preferences over all higher moments of the point-in-time distribution of returns. In a mean–variance world where all asset returns are Gaussian, or where the investor has preferences only over the first two moments, this reduces to σ^2 as a special case.

σ_B^2 is behavioural variance, a generalisation of variance that incorporates the investor’s preferences for all the higher moments of the distribution for investors with constant relative risk aversion (CRRA) utility for log returns on total wealth. With mean–variance preferences, or when all returns distributions are Gaussian, the measure collapses to variance as a special case.

However, anxiety will induce interim decisions that cause the investor to deviate from the normatively optimal portfolio (portfolios with too much anxiety along the journey will be unattainable in practice). The rational investor will therefore take account of their own behavioural distortions and seek to reduce anxiety by choosing normatively suboptimal portfolios, which nonetheless offer the best attainable result. The best realisable expected utility is therefore a function of both desirability and anxiety, increasing in the former and decreasing in the latter. It will be difficult to establish the precise nature of this trade-off. Anxiety, by its nature, is likely to be highly unstable and idiosyncratic, and may stem from the interaction of two distinct sets of preferences: one stable, long-term, normative and financial (desirability); one unstable, immediate, behavioural and emotional (anxiety). In this context it may be meaningless to speak of a single, coherent preference ordering that would govern the optimal trade-off between desirability and anxiety. Preferences would be constructed afresh in each context and situation (Slovic 1991). However, without requiring us to specify a stable trade-off function, we can borrow a technique familiar to portfolio theory, the “efficient frontier”.

While we may be unable to say what the optimal trade-off is, as long as anxiety is reliably generated by certain characteristics of the possible paths of the investment journey, we can at least establish the

set of portfolios that constitute an efficient trade-off between emotional comfort and normative financial efficiency. In other words, whatever degree of comfort is required for the investor to stick with their choices through time, we can determine the portfolios that achieve this at minimum cost.

DEFINING ANXIETY

To do this we need to have some proxy measure for the degree to which different portfolio choices might produce anxiety, analogous to using a risk measure to represent the expected-utility-reducing aspects of a portfolio in satisfying our purely financial preferences.

As we see from the discussion above, a wide range of factors might influence the comfort the investor feels with their portfolio at any point in time. Some of these are highly idiosyncratic and specific to the investor. Behavioural responses may be prompted by specific events that are not a general feature of the portfolio itself: a passing comment by a friend, perhaps; happening to observe an asset at a particular time; or general unease from other aspects of the investor's life that translate into their discomfort with normative portfolio decisions. We cannot hope to systematically include all idiosyncratic behavioural influences into a measure of anxiety. But we can seek to examine characteristics of portfolios that are more or less likely to induce anxiety (ie, reduce comfort with the normative solution) along the journey.

In what follows we shall leave aside the question of experiences that result in investors distorting probability distributions, although we acknowledge that this aspect certainly might cause deviations from a normative solution. For instance, a recent period of bad returns may induce investors to believe negative events are more likely than their previous assumptions had led them to believe, thereby placing greater weight in the left tail of the distribution and reducing subjective expected returns. Probability distortion may certainly cause changes in investors' evaluations of portfolios, but here we assume that they retain their assumptions on the multivariate distribution of outcomes (perhaps because they rely on quantitative models or advisors to inform their expectations of the future) and instead only change their preferences for those outcomes. This will

enable us to examine the effects on rational portfolio choice of changing short-term preferences without also having to model changing beliefs.

Anxiety from distorted point-in-time utility

Davies and de Servigny (2012, Chapter 8) take one approach to this that is in the spirit of prospect theory. They model short-term behavioural preferences as psychophysical distortions of the same distribution of potential future returns as the normative financial preferences. That is, when considering the multivariate distribution of return at the rebalancing horizon, two sets of preferences sequentially come into play. First, the short-term behavioural or psychophysical distortions (reflecting the first set of non-financial, reference-dependent preferences) transform the returns, before these returns are entered as inputs into the second set of normative expected utility preferences. Thus, normative preferences still establish the optimal portfolio, but perceptions of returns are behaviourally, and unavoidably, distorted prior to the calculation of expected returns.

This approach contains both normative and behavioural preferences in a single combined function, applied to the point-in-time probability distribution. For a fixed set of parameters this permits the explicit incorporation of the anxiety measure,⁸ A , into a revised expected utility calculation, such that

$$D = (\bar{r} - r_f) - \frac{\sigma_B^2}{T} - A$$

But it fails to capture what we consider to be an essential feature: anxiety is much more about the experience of the journey than about the distribution of outcomes at the destination.

If the planned rebalancing point is a long way off, say annually, then there is ample scope for the investor to respond behaviourally to features of the path along the way, thereby deviating from the normative portfolio well before the reassessment point arrives. Protecting against anxiety measured on the point-in-time distribution 12 months hence will miss emotionally important aspects of the journey.

This is less of an issue if the reassessment point is frequent, so there is not much that can happen along the path that does not show up in the end distribution. However, such near point-in-time distributions

also cannot be expected to reflect the full effects of anxiety. Stress and discomfort are unlikely to build only between one reassessment point and the next, with all behavioural distortion swept clear as soon as each checkpoint is reached. Instead, it may build cumulatively over time until we run out of our reserves of emotional liquidity and give in to our craving for emotional comfort.

For example, if markets have been falling, then investors feel level returns as a loss (the reference point is above 0%); or if markets have been rising strongly, then investors face the house money effect where subsequent losses are experienced merely as reduced gains (the reference point is below 0%). Davies and de Servigny do consider this possibility and show how it may be incorporated into their approach by varying the position of the reference point, but they do not provide any model for how to determine the likely position of the reference point given the realised past history, so do not develop a practical solution to this issue.

Anxiety over the journey

To resolve these concerns we propose extracting a measure of anxiety from the paths themselves, thereby making anxiety explicitly about the experience of the journey, as distinct from desirability, which is a characteristic for the point-in-time outcome distribution. What we seek is a method to determine the degree to which any given path is likely to make investors uncomfortable with the investment they are in, such that they are tempted to deviate from it in order to make themselves more comfortable. We shall examine a number of characteristics of paths that behavioural research indicates investors are likely to find uncomfortable and, while we cannot claim to have arrived at some “optimal” measure of anxiety, we choose a single plausible and tractable candidate to investigate the efficient desirability–anxiety (D–A) trade-offs that a rational, self-aware investor might make.⁹

Path volatility as anxiety

An obvious first candidate for measuring the experience of an investment journey is one frequently confused with risk. Volatility – the standard deviation of period-by-period returns along the path – is essentially a measure of the “roughness” of the journey. It is quite distinct from risk, which is a characteristic of the distribution of outcomes at the destination, though the latter is often also measured

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

using the standard deviation of these outcomes, and rough journeys frequently correlate significantly with the risk of a poor outcome.

In measuring anxiety, we are only concerned with a poor outcome as a consequence of investors' actions in response to the emotional stress of a poor journey. Quite plausibly, a path with higher standard deviation is one with lower emotional comfort, and therefore one with a greater likelihood of costly behavioural responses along the way.

But is volatility a good measure of the anxiety of a returns path? It is certainly more stressful to experience a volatile journey than a smooth one, but volatility accounts neither for the trend nor for salient points along the way (peaks or troughs, for example). We can imagine a highly volatile path that is nonetheless strongly increasing (so the investor feels relatively comfortable, despite the turbulence), or a very smooth path that is steadily falling (thus causing anxiety despite very low volatility). Volatility treats positive and negative deviations from the trend symmetrically, whereas evidence on reference dependence and loss aversion, and our intuition, would indicate that downward movements cause far more stress than increases. It also does not account for the order in which returns happen, which would seem to be a crucial consideration when evaluating temporally experienced returns: a path with six negative months in a row followed by six positive months would to many to be intuitively more stressful than a path with the same twelve returns reordered so that up and down months alternate sequentially. They have the same volatility, but the first has the stressful cumulative effect of six losses in a row. So, while taking steps to reduce volatility will reduce anxiety, the relationship is not as close as we would like, meaning that sacrifices of expected utility in pursuit of volatility reduction will purchase less comfort than if we had a more targeted measure.

Anxiety-generating path characteristics

To overcome these concerns it will help to ask what features of investment paths will contribute most to emotional discomfort along the journey, and then seek a measure that captures these.

First, we would want our measure to place more emphasis on losses than gains. This is not to say that gains do not also induce behavioural responses: tendencies for exuberance, trend-following and herd behaviour are all examples of unexpectedly good investment returns leading to temptations to deviate from the financially

optimal portfolio. However, loss aversion implies that decreases will have much greater emotional impact than increases, and as a general rule myopia will incline most investors towards taking considerably less risk than is appropriate for their longer-term financial preferences. As a result, for most investors exuberance will largely serve a corrective function to the more usual behavioural effects of short-term nervousness.

Second, experimental explorations show that retrospective evaluations of past experiences can be largely predicted by averaging just two aspects of the path – how individuals felt at the most intense point along the path, and how they felt at the end (the “peak–end rule”) – while the duration of the experience was relatively unimportant (“duration neglect”).¹⁰ We would therefore expect the most salient moments of the path, the peaks and low points of portfolio value, to have disproportionate effect on overall anxiety, and the intensity or magnitude of the loss (or gain) to lead to greater discomfort than the length of time over which it occurs.

Third, memory decay and our ability to adapt to new circumstances (the “treadmill” effects discussed in Kahneman 2000) would imply that the discomfort experienced along a particular path does not permanently raise anxiety. Aspects of a path that are in the distant past will not continue to affect anxiety levels for that portfolio indefinitely.¹¹ For this reason we would wish to impose a statute of limitations on the influence of past events on the anxiety measure.

Finally, we would wish to account for adaptation to the path itself (ie, how individuals’ reference points update in response to the path), which then influences whether outcomes are perceived as gains or losses. We may well expect faster reference point updating in response to increases than when wealth declines. Evidence that people prefer steadily increasing paths for many aspects of their life (Loewenstein and Prelec 1993) indicates that people can fairly rapidly become acclimatised to improving fortunes,¹² but that their expectations are fairly sticky on the way down: we tend to cling to past glories and find it more difficult to adapt to declining fortunes.

Drawdown as a stylised measure of anxiety

While we make no claim that it is the “best” measure of path anxiety, there is one measure that embodies almost all of these characteristics, and that has the added benefit that it is already widely used and

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

understood in the finance industry: drawdown – the magnitude of the fall from a peak value to the lowest subsequent point on the path. It is very much a measure of a temporal path, rather than a point-in-time distribution. It is inherently focused on losses, rather than gains, thus meeting our first requirement above. Indeed, it may be too extreme in this regard, since gains along the path only feature insofar as they set the peak point.¹³ Drawdown is also driven by two of the most salient, and extreme, points of any path: the peak and the point of largest loss from that peak. It is extremely intuitive that this drop is a feature of the journey that matters a great deal to investors: the fact that drawdown is so widely used and quoted despite not playing any formal role in classical portfolio theory is testament to its perceived importance. Next, the measure is purely about the magnitude of the drop, and not about its duration: a rapid drop of 20% has the same drawdown value as the same drop over several months. This matches evidence on duration neglect. This is not to say that duration does not matter at all, but rather that in drawdown we have a relatively simple, pared-down measure that captures much of what is most important about the path, while leaving aside lesser aspects that might have some role to play but would represent a poor increase in accuracy for the increased complexity they would bring.

The only one of the above requirements drawdown does not naturally incorporate is the third point, on not including all past events. Drawdown is not typically defined over a specific period, so would include all past peaks, no matter how far back they are. In truth this is a drawback to the measure, since it makes drawdown figures difficult to compare consistently. However, imposing a specific time period over which to measure drawdown easily solves this. This is a somewhat arbitrary choice, but we favour looking at drawdown over a constant 12-month period, largely because the period of a year has enormous social and financial import. We all face calendar years, tax years, reporting on an annual basis and annual salary reviews. So if we are to choose a single window within which to examine anxiety, it makes sense to match the window that most people will find most natural as the length of time over which they keep a number of important mental accounts open.

The last point is reference point updating, with regard to which drawdown has a particularly interesting psychological interpretation. Drawdown may be seen as a measure of the peak negative experience in a reference-dependent model with an extreme version of dynamic, path-dependent updating of the reference point. Essentially, it assumes that reference points update instantly to gains (ruling out any possibility of the house money effect), but are infinitely sticky with regard to losses. So, for each path, the best point sets an investor's reference point, and the peak of anxiety is the total loss relative to this point. While this is extreme, it nonetheless reflects, in a simple and stylised way, the intuition that we are more adaptable to gains than to losses. We could expect that in reality a reference point would not update instantly on the way up, leaving some emotional buffer against subsequent losses, and the reference would show some adjustment on the way down as investors gradually acclimatise to their new situation.¹⁴ But this would greatly complicate the measure, and also make it less readily understood and communicated to users.

So, while there are many tweaks we might wish to consider, particularly a more sophisticated model of path-dependent reference point updating, drawdown nonetheless captures in a simple and familiar way all the essential features that we would wish to see of a measure of anxiety. In what follows we shall examine the nature of the trade-off between expected return and anxiety, using 12-month drawdown as our measure. We shall also examine a number of approaches to the efficient reduction of anxiety: purchasing emotional comfort at the lowest cost to our normative financial preferences.

AN EMPIRICAL EXPLORATION OF ANXIETY REDUCTION AT LOWEST COST

To demonstrate this notion we now undertake an illustrative exploration of some possible mechanisms for efficiently reducing anxiety. These split into three broad categories:

1. altering investor experiences of paths by deviating from the normative asset allocation;
2. altering investor perceptions of the journey with no changes to the portfolio; and

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

3. modifying the paths directly, through the use of structuring or derivatives.

In changing the asset allocation we show how a reduction in anxiety may be acquired efficiently (ie, with limited reduction in expected utility from the perspective of financial preferences) by judicious choice of alternative asset allocations. This is less costly than (but not as simple as) investors' natural inclination to reduce anxiety through underinvestment, thereby mixing the normative solution with cash. We then assess the effect of altering investor perceptions by reducing the frequency at which the investor monitors the portfolio. This is the cheapest method (by virtue of being free) but requires consistent discipline if it is to be successful. Lastly, we discuss how modifying the realised path of the portfolio through the use of derivatives and structuring affects both desirability and anxiety together and, depending on option pricing, may offer an efficient alternative to changing the asset allocation alone.

We start with a set of three asset portfolios comprised only of cash, bonds and equities, which we represent by three-month US Libor, Barclays Global Aggregate Treasuries USD hedged, and MSCI World Net (USD),¹⁵ respectively. We construct all possible long-only portfolios combining these assets, where the weights are multiples of 2%. For each we examine two possible historical periods of monthly total return data: the 1990s and the 2000s. Of course when optimising a portfolio we want to consider the future returns distribution: here we simply use the historical data from these two periods as proxies for the sort of risk–return characteristics we may have for portfolios of these three assets.

The normatively optimal solutions

Figures 5.1 and 5.2 show the conventional risk–return characteristics for the two data periods, where each point represents a single portfolio's monthly mean total return and behavioural standard deviation. An investor facing a choice between these portfolios, and concerned solely with optimising financial preferences for risk and return over the next month (when they can reassess and shift to a new optimal portfolio), would pick one of the points on the risk–return efficient frontier.

DON'T LET THE BEST BE THE ENEMY OF THE ACHIEVABLE

Figure 5.1 Monthly excess return (over risk-free rate) versus behavioural risk, January 31, 1990, to December 31, 1999

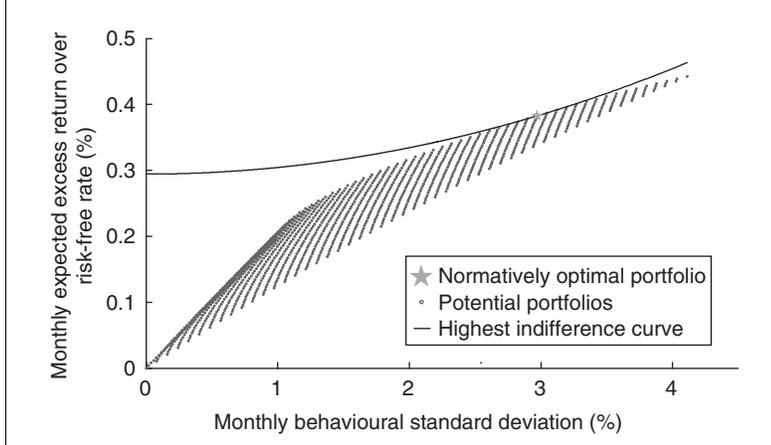
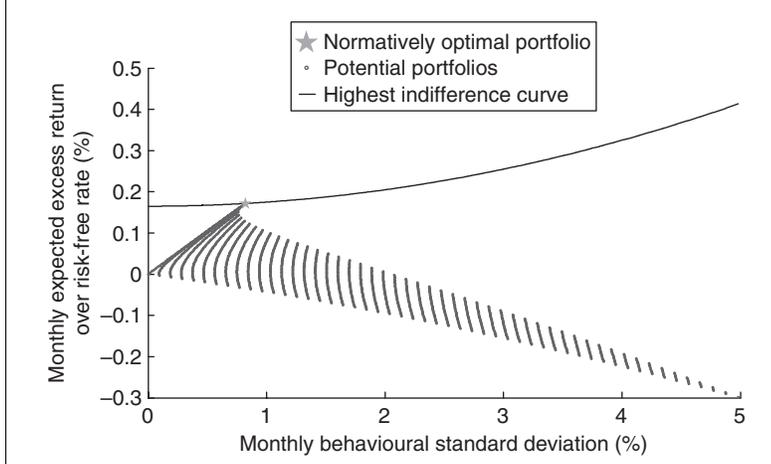


Figure 5.2 Monthly excess return (over risk-free rate) versus behavioural risk, January 31, 2000, to December 31, 2009



The point on the frontier with the greatest expected utility is the one that has the highest desirability, given the investor's risk tolerance, T . This value is the intercept on the y -axis of the highest indifference curve that can be achieved for this set of portfolios. Throughout this analysis we shall look through the eyes of a single investor, with $T = 1$, for whom the optimal choice on this curve

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

is at the point of tangency between the frontier and the investor's indifference curve.¹⁶

You will note that the efficient frontier in the 2000s consists only of mixes between cash and bonds: it was a period when greater risk was not remunerated. On an *ex ante* basis we would not expect this to be the case: models such as Black and Litterman (1992) establish forward-looking equilibrium expected returns that assume assets earn a premium commensurate with their risk. We show this period for completeness, and the existence of such periods, or even of investors holding such beliefs for the future, does not alter our fundamental principle of needing to efficiently reduce anxiety. Where such expectations do exist, however, they make anxiety reduction more difficult. In the analysis that follows we shall predominantly use the data from the 1990s as a proxy for the sort of reasonable expectations of future returns that investors typically face when determining the optimal portfolio.

The normative solutions shown in Figures 5.1 and 5.2 are with reference to only two dimensions: risk and return. The optimal portfolio (picked out by a "star") is the most efficient for that investor with the degree of risk tolerance indicated by the indifference curve. Desirability provides a single summary measure of the efficiency of combinations of these two dimensions: a single number thus describes portfolio efficiency in the normative world. The difference in desirability between the "star" and any other portfolio shows the loss in efficiency from choosing the latter. In the following figures, instead of retaining separation between these two dimensions, we use the combined measure of desirability to reflect normative efficiency.

Optimal achievable solutions: efficient anxiety reduction

The portfolio with the highest desirability is the normatively optimal portfolio for each investor but, from our discussion above, may not be the optimal/rational portfolio for the investor to choose in practice. Any investor with sufficient self-awareness will be concerned not just with the theoretical optimum but also with their ability to enact this portfolio, and stick with it. Our investor should therefore seek to avoid any portfolios that have particularly high expected anxiety along the journey, and be prepared to sacrifice some financial efficiency to secure a higher achievable return. If an investor is going to forgo portfolio efficiency (ie, desirability) to reach their

required level of anxiety reduction, they should rationally sacrifice as little as possible.

Because of both the highly unstable nature of behavioural distortions (and therefore of investors’ susceptibility to anxiety) and the fact that these preferences are largely distinct from financial preferences (anxiety derives from the path, not the point-in-time distribution), we do not believe it is possible to coherently combine preferences for desirability and anxiety into a single “meta” utility function. That is, we cannot construct a stable indifference curve to determine the single “optimal” combination of anxiety and desirability. However, we can say that, in some sense, satisfaction of overall preferences is increasing in the former and decreasing in the latter. So while we cannot pick an “optimal” portfolio, we can, in a similar vein to the traditional risk–return efficient frontier, define a desirability–anxiety (D–A) efficient frontier.

As discussed above we need to ensure that our anxiety measure does not have infinite memory. In keeping with the ubiquitous focus on one year as the most socially prevalent time horizon, we confine our proxy for the characteristic (non-investor specific) anxiety of each path, drawdown, to a 12-month period. This 12-month drawdown, DD, is given as

$$DD = \max_{0 \leq t \leq 12; t \leq n \leq 12} \left(\frac{v_t - v_n}{v_t} \right)$$

where v_t is the value of (or a proxy for the perceived value of) the portfolio at time t . We use a bootstrap-with-replacement method to produce 2,000 test paths over time. To determine a measure of expected anxiety for the portfolio as a whole we take the mean maximum drawdown over all portfolio paths, denoted by $\mathbb{E}[DD]$. In generating the paths we assume portfolios rebalancing monthly, not buy-and-hold.

Figures 5.3 and 5.4 show all portfolios in desirability– $\mathbb{E}[DD]$ space, for each decade. Concentrating on Figure 5.3, the highest point, highlighted with a star, would be the normatively efficient portfolio to pick. Although this point has the highest desirability, the associated expected drawdown¹⁷ may be more than sufficient to cause reluctance or adverse behavioural responses from even slightly nervous investors. Classical finance nonetheless maintains that this is the rational portfolio to choose, and the level of anxiety plays no part in the decision.

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

Figure 5.3 Plot of $\mathbb{E}[DD]$ and desirability highlighting desirability–anxiety frontier portfolios, January 31, 1990, to December 31, 1999

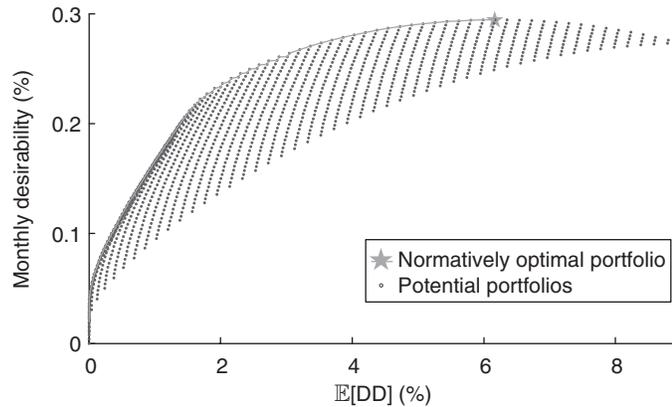
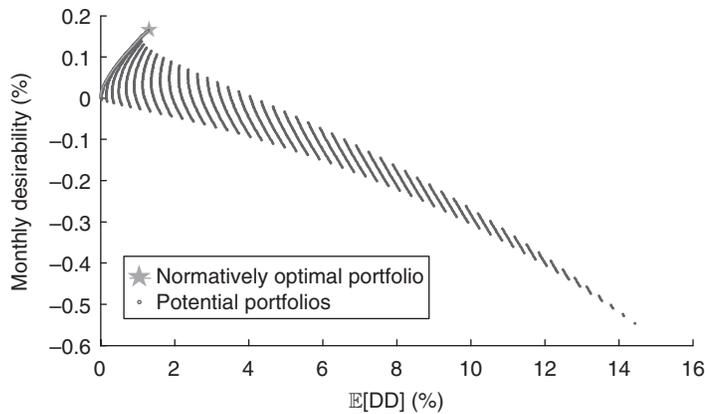


Figure 5.4 Plot of $\mathbb{E}[DD]$ and desirability highlighting desirability–anxiety frontier portfolios, January 31, 2000, to December 31, 2009



All other portfolios result in the investor sacrificing some expected utility of the outcome, but the portfolios on the line forming the outer hull of the set further left of this point and downwards are those that make this sacrifice most effective. Particularly notable is the first part of this curve, immediately to the left of the normative choice: these portfolios show very rapidly declining $\mathbb{E}[DD]$, in

exchange for small reductions in financial efficiency.¹⁸ The difference between our two time periods makes it clear that the rate at which we could decrease $\mathbb{E}[DD]$ by forgoing desirability is dependent on the anticipated environment, or equivalently, our beliefs about the future.

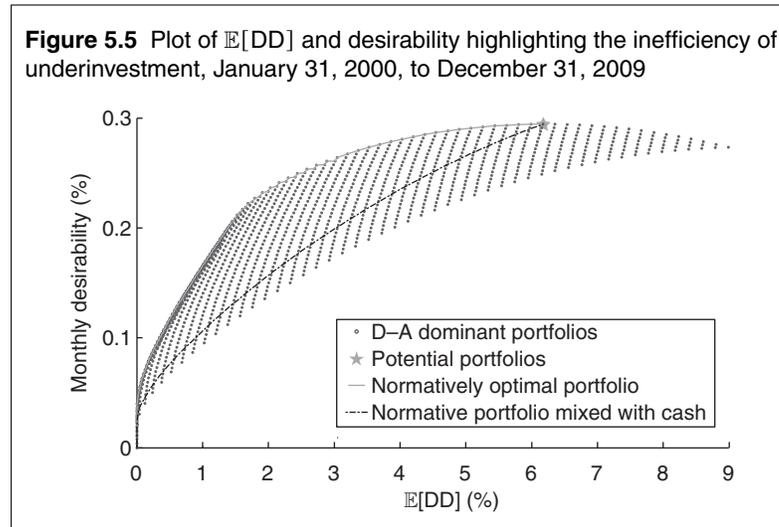
We cannot say which portfolio embodies the optimal trade-off, but can posit that the more nervous or present-focused the investor, the further down and left the portfolio they should ideally settle on. The investor may, for example, ascertain that any portfolio with $\mathbb{E}[DD]$ greater than a certain amount would be dangerously uncomfortable, and therefore settle on the portfolio with the greatest expected utility that has $\mathbb{E}[DD]$ below this level, effectively the solution to a constrained optimisation.

It is worth noting too that all the D–A efficient portfolios in this environment are also risk–return efficient (or indistinguishable from being so): by reducing anxiety, the investor is effectively moving to the left along the traditional (expected return versus behavioural risk) efficient frontier. Importantly, this held for both sets of extremely different data. This need not always be the case (if, for example, rather than just picking portfolios, we were to use derivatives to actively transform the paths themselves), but it is compelling that in both these different cases we can enhance comfort efficiently using the same set of portfolios already picked out by traditional risk–return optimisation (although we need to deviate from the normative portfolio).

Comparison to underinvestment as a means to reduce anxiety

Anxiety-induced fear of investment often leads to investors acquiring comfort in the simplest, most intuitive manner, by underinvestment, typically retaining a large holding in cash-like assets. This is a potentially suboptimal way of self-medicating investment angst. To explore how our D–A efficient portfolios compare with this simple strategy, we examine the costs to expected utility of investors choosing the normatively optimal portfolio, and then failing to invest fully; they thus hold a mix of the highest desirability portfolio and cash. The dot-dashed line in Figure 5.5 shows this. The solid line in Figure 5.5 is formed of the “D–A dominant portfolios”, ie, those portfolios that have highest desirability for a given $\mathbb{E}[DD]$.

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT



By definition, this solid line of portfolios will always be at least as high as (eg, on the 2000s data), and usually higher than, the dot-dashed line showing mixes of the normative portfolio with cash: hence, underinvestment is usually an expensive way of reducing $\mathbb{E}[DD]$. The unnecessary cost of underinvestment is shown by the distance between the two (solid and dot-dashed) lines in Figure 5.5.

Thus, this mix of cash with the “star” portfolio turns out to be a highly inefficient way of reducing $\mathbb{E}[DD]$ and, depending on the forward-looking assumptions, could be extremely costly in the longer term.¹⁹ Again, it is exciting that in order to improve investor D–A efficiency from the natural response of underinvestment we do not have to do anything more complex than select a more conservative point on the conventional risk–return frontier.

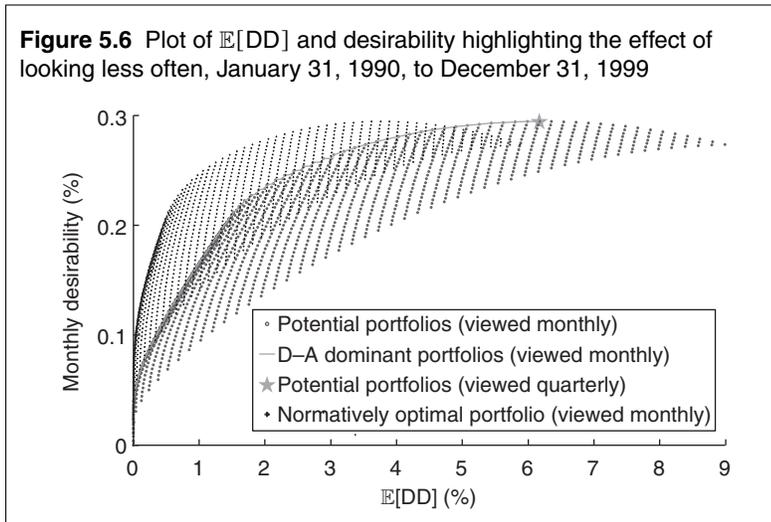
Underinvestment is a common occurrence, and in some cases may be an unconscious effort towards anxiety reduction, but with it the investor gives up more desirability than needed.

Targeting anxiety through perceptions of the path

Before we accept these desirability trade-offs we need to examine whether there are cheaper ways to achieve this.

Figure 5.6 depicts the same portfolios (still rebalanced monthly) but $\mathbb{E}[DD]$ has been recalculated for all portfolios, assuming the investor only looks at or notes the investment returns every quarter. As monthly drawdown peaks and troughs will not always coincide

DON'T LET THE BEST BE THE ENEMY OF THE ACHIEVABLE



with the end of a quarter, and this reduces the quarterly drawdown, the result is that the anxiety measure for all portfolios shifts to the left. Of course, the underlying portfolios are no different, and the actual path they take does not change, but this effect reveals a crucially important aspect of anxiety: it is as much about perception as reality. As a result of less frequent observation, the intensity of experienced discomfort along the path decreases, and the D–A efficient frontier appears to shift left. This is the cheapest (with regard to expected utility) reduction of anxiety, as there is no effect on the outcome distribution.²⁰

While reduced observation seems to reduce anxiety for free, and is certainly a recommended strategy for many investors of nervous disposition, there may be practical limits to how far this can be pushed. For example, with monthly rebalancing, quarterly observation means paying someone to rebalance for you, which is costly; and in times of crisis it may be virtually impossible not to note that asset prices are plummeting.

This type of self-control is usually the cheapest method of anxiety reduction but only if the investor can adhere to it.

Targeting anxiety by controlling the path

So far we have considered three simple methods of attempting anxiety reduction:

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

- investing in a portfolio optimised to reduce anxiety directly through cost to desirability;
- inefficient, costly but simple mixing of the normative portfolio with cash;
- less frequent monitoring and no change to the portfolio itself.

An additional method to target anxiety reduction could involve changing assets in order to manipulate the paths themselves. This will also change the shape of the end distributions. There are many ways to influence the paths of an asset using derivatives, for example, though buying a put option, which places a floor under the final value of the investment. The cost of this protection reduces the final returns of all outcomes, and shifts all final returns below the floor to the floor (minus the option cost): this changes the final returns distribution, and therefore desirability. But it also enhances the investor's comfort relative to the underlying asset whenever it dips below this floor along the journey, thereby reducing anxiety. As long as the effect on anxiety is greater than the effect on desirability, this could provide a more efficient way of reducing anxiety.

If options are sufficiently cheap and have sufficient effect on the path, then portfolios could be re-optimised in a way that may make investors more comfortable with their investment path at a lower cost to financial efficiency than choosing among the standard portfolios. The pertinent word is sufficient: an option affects both the path and point-in-time desirability, and to be an efficient means of anxiety reduction the former needs to outweigh the latter.

Collars and other such varieties of options could offer more cost effective anxiety reducers but how worthwhile they are would depend on the premium, the anticipated multivariate returns distribution and the investor's risk tolerance (due to its influence on desirability).

Although we have not had the scope to do so here, we believe that exploring derivative-based methods of altering the journey itself would be a valuable direction for further work.

CONCLUSION

We have not by any means explored all possible methods for efficiently reducing anxiety. For example, while tactical asset allocation is conventionally used to pursue greater risk-adjusted returns or

DON'T LET THE BEST BE THE ENEMY OF THE ACHIEVABLE

greater desirability, dynamic portfolio allocation could instead be deployed to serve the needs of anxious investors, by targeting anxiety reduction. However, we have illustrated simple demonstrations of the central concept: that attaining the best achievable returns, given our behavioural responses to the journey, may involve deliberately turning away from the normatively optimal portfolio to get sufficient comfort to stick with the journey.

Our simple analysis indicates that the best value anxiety reduction comes from the investor altering their perception of paths where possible, rather than deviating from the normative portfolio solution, but in all likelihood some combination of all these approaches may be the most beneficial in reality.

In summary, rationality encompasses more than a normative solution. It also requires that an individual is able to implement that strategy from the start, and stick to it through to the finish. The realised path of the investment journey plays an important part in our ability to govern our emotional responses, and a rational solution takes this and our own behavioural fallibility into account. In this chapter we have advanced some initial ideas of how this solution might be achieved: from altering perception to changing the investment itself. Of course, the real test is in the emotional responses of investors: was the journey more comfortable and, ultimately, were they able to stick to it and thereby attain better returns?

Portfolio theory should not ignore the primary reason that investors fail to attain good returns, that is, not a lack of knowledge of the normative solution, but the lack of the emotional capacity to enact it. Instead, portfolio theory should expand its remit to provide investors with the best possible attainable solution by embracing financially efficient anxiety reduction. Do not let the normative be the enemy of the achievable.

- 1 The choice of risk measure used to construct the frontier also implies assumptions about preferences. A mean–variance frontier only picks out optimal combinations under the assumption of quadratic utility, or when asset returns are Gaussian.
- 2 Except insofar as they provide information that allows us to update our expectations of the future probability distribution.
- 3 Baddeley (2013) is a good starting point.
- 4 This is not to say that this is always the case: in times of irrational exuberance a combination of diminished risk perception, herd behaviour and anticipated regret of falling behind others in their investing success will mean that the behaviourally comfortable route is excessive risk taking, rather than too little.
- 5 The result of buying high and selling low (Dichev 2007; Friesen and Sapp 2007; Dalbar 2013).

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

- 6 We need to distinguish between this time horizon (at which we assess the distribution) and the horizon at which the investor needs the wealth to fund future consumption or meet varied future expenditure goals. The latter will partly determine the appropriate degree of risk tolerance, T , in the utility function. The shorter the weighted expected duration of the cash outflows, the lower risk capacity, effectively lowering T .
- 7 They make the case for a normative model where T governs the curvature of CRRA utility on log returns of total wealth.
- 8 Though Davies and de Servigny caution that A is inherently unstable and should be used to make broad comparisons between investments, not to pick a precise portfolio optimising the combined preferences.
- 9 A word of caution about the terms "emotional comfort" and "anxiety": there could be times when the problem is not too little comfort with investing but too much. For example, facing the fear of "missing out" and falling behind friends and colleagues after markets have been rising for some time, and therefore increasing risk exposure at the top. This is still a behavioural deviation from the normative solution, albeit opposite to the more common situation of too little risk taking due to myopic loss aversion. By anxiety and emotional discomfort we mean discomfort with the normative solution, regardless of the direction of this discomfort.
- 10 See Fredrickson and Kahneman (1993) and Kahneman (2000).
- 11 An exception to this may be events imbued with particular personal import for the investor: for example, losses incurred due to a decision that was uncharacteristic for the investor, and which they therefore regretted; or the purchase price of an investment, which may have psychological salience for much longer than interim experiences. The vividness of events such as these could dramatically increase their influence on the investor's reference point, and thus on whether certain outcomes are perceived as gains or losses. However, while important, these events are idiosyncratic to the investor and therefore cannot be effectively incorporated into a general measure of path anxiety.
- 12 Though the existence of the "house money effect" would imply that there are certainly circumstances where we do not adapt fully and instantaneously to large, prolonged or rapid increases in wealth.
- 13 This could be resolved by combining drawdown in some way with a measure of "drawup", which would reflect the magnitude of subsequent increases from the lowest point of a path. The relative importance of losses over gains would require that drawup has a lower weight than drawdown in any combined measure.
- 14 This is accounted for to some degree by the 12-month memory we have imposed.
- 15 Total return data sourced from Factset and Bloomberg. Equivalent Bloomberg tickers: US0003M Index, LGTRTRUH Index, NDDUWI Index.
- 16 We make this assumption for consistency and ease of comparison. Different levels of risk tolerance will change the desirability numbers for all portfolios, and thus mean a different optimal portfolio, but will not change the general nature of the desirability–anxiety trade-off we are illustrating. $T = 1$ is perhaps slightly more risk tolerant than a typical investor.
- 17 We make no claim for these numbers to have general significance outside of the chosen data set, but they give some indication of the scale of $\mathbb{E}[\text{DD}]$ relative to desirability. In this example, the desirability is 0.3% (or 3.5% annualised), and the expected 12-month drawdown $\mathbb{E}[\text{DD}]$ is 6.2%.
- 18 For example, on the 1990s data set, the investor could reduce $\mathbb{E}[\text{DD}]$ from just over 6% to 4% while barely sacrificing any financial efficiency (annualised desirability reduces from 3.5% to 3.3%). Using the 2000s data, the most desirable point already has $\mathbb{E}[\text{DD}]$ far below 4%, and any further reduction is more costly.
- 19 Were our forward-looking beliefs to be based purely on our historical experience of the 2000s, however, this would not be the case, since all risk–return efficient portfolios are a mix of cash and the optimal portfolio of 100% bonds, and the optimal D–A line and cash-mix line coincide.
- 20 This effect of anxiety decreasing in observation frequency is also found with the point-in-time anxiety measure of Davies and de Servigny (2012).

REFERENCES

- Baddeley, M.**, 2013, *Behavioural Economics and Finance*, Routledge Advanced Texts in Economics and Finance (Oxford: Routledge).
- Barber, B. M., and T. Odean**, 2000, “Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors”, *Journal of Finance* 60, pp. 773–806.
- Barberis, N., and W. Xiong**, 2011, “Realization Utility”, *Journal of Financial Economics* 104(2), pp. 251–71.
- Benartzi, S., and R. H. Thaler**, 1995, “Myopic Loss Aversion and the Equity Premium Puzzle”, *Quarterly Journal of Economics* 110, pp. 73–92.
- Black, F. and R. Litterman**, 1992, “Global Portfolio Optimization”, *Financial Analysts Journal* 48(5), pp. 28–43.
- Dalbar**, 2013, “Quantitative Analysis of Investor Behaviour”, Report, Research and Communications Division, Dalbar Inc.
- Davies, G. B., and A. de Servigny**, 2012, *Behavioral Investment Management: An Efficient Alternative to Modern Portfolio Theory* (New York: McGraw-Hill).
- Dichev, I. D.**, 2007, “What Are Stock Investors’ Actual Historical Returns? Evidence from Dollar-Weighted Returns”, *American Economic Review* 97, pp. 386–401.
- Dorn, D., and P. Sengmueller**, 2009, “Trading as Entertainment?”, *Management Science* 55(4), pp. 591–603.
- Frazzini, A., and O. A. Lamont**, 2008, “Dumb Money: Mutual Fund Flows and the Cross-Section of Stock Returns”, *Journal of Financial Economics* 88, pp. 299–322.
- Frederickson, B. L., and D. Kahneman**, 1993, “Duration Neglect in Retrospective Evaluation of Affective Episodes”, *Journal of Personality and Social Psychology* 65, pp. 45–55.
- French, K. R., and J. M. Poterba**, 1991, “Investor Diversification and International Equity Markets”, *American Economic Review* 81, pp. 222–6.
- Friesen, G. C., and T. R. A. Sapp**, 2007, “Mutual Fund Flows and Investor Returns: An Empirical Examination of Fund Investor Timing Ability”, *Journal of Banking and Finance* 31, pp. 2796–816.
- Frydman, C., N. Barberis, C. Camerer, P. Bossaerts and A. Rangel**, 2014, “Using Neural Data to Test a Theory of Investor Behavior: An Application to Realization Utility”, *Journal of Finance*, URL: <http://dx.doi.org/10.1111/jofi.12126>.
- Genesove, D., and C. Mayer**, 2001, “Loss-Aversion and Seller Behavior: Evidence from the Housing Market”, *Quarterly Journal of Economics* 116(4), pp. 1233–60.
- Ingersoll, J. E., and L. J. Jin**, 2013, “Realization Utility with Reference-Dependent Preferences”, *Review of Financial Studies* 26(3), pp. 723–67.
- Kahneman, D.**, 2000, “Experienced Utility and Objective Happiness: A Moment Based Approach”, in D. Kahneman and A. Tversky (eds), *Choices, Values, and Frames* (Cambridge University Press).
- Kahneman, D.**, 2011, *Thinking, Fast and Slow* (London: Allen Lane).
- Kahneman, D., and A. Tversky**, 1979, “Prospect Theory: An Analysis of Decision under Risk”, *Econometrica* 47, pp. 263–91.
- Kumar, A.**, 2009, “Who Gambles in the Stock Market?”, *Journal of Finance* 64(4), pp. 1889–1933.

QUANTITATIVE APPROACHES TO HIGH NET WORTH INVESTMENT

Loewenstein, G. F., 1993, "Preferences for Sequences of Outcomes", *Psychological Review* 100, pp. 91–108.

Loewenstein, G. F., and D. Prelec, 1993, "Preferences for Sequences of Outcomes", *Psychological Review* 100(1), pp. 91–108.

Loomes, G., and R. Sugden, 1982, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty", *Economic Journal* 92, pp. 805–28.

McNeil, A. J., R. Frey and P. Embrechts, 2005, *Quantitative Risk Management: Concepts, Techniques and Tools* (Princeton University Press).

Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", *Journal of Economic Theory* 3(4), pp. 373–413.

Odean, T., 1998, "Are Investors Reluctant to Realize Their Losses?", *Journal of Finance* 53, pp. 1775–98.

Rabin, M., 2000, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem", *Econometrica* 68, pp. 1281–92.

Samuelson, P. A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming", *Review of Economics and Statistics* 51, pp. 239–46.

Sloman, S. A., 2002, "Two Systems of Reasoning", in T. Gilovich, D. Griffin and D. Kahneman (eds), *Heuristics and Biases: The Psychology of Intuitive Judgment* (Cambridge University Press).

Slovic, P., 1991, "The Construction of Preference", *American Psychologist* 50, pp. 364–71.

Statman, M., 2010, *What Investors Really Want: Know What Drives Investor Behavior and Make Smarter Financial Decisions* (McGraw-Hill).

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

Thaler, R. H., and E. J. Johnson, 1991, "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice", in R. H. Thaler (ed), *Quasi-Rational Economics* (New York: Russell Sage Foundation).

Tversky, A., and D. Kahneman, 1992, "Advances in Prospect Theory: Cumulative Representation of Uncertainty", *Journal of Risk and Uncertainty* 5, pp. 297–323.

Weber, M., E. U. Weber and A. Nasic, 2012, "Who Takes Risks When and Why: Determinants of Changes in Investor Risk Taking", *Review of Finance* 17, pp. 847–83.