

Understanding goal-based investing

By Joao Frasco, Chief Investment Officer, STANLIB Multi-Manager



This article will explain our thinking behind goal-based investing.

It is important to understand that we are only talking about the investment side of meeting goals here, and that financial planning is a much broader topic that will directly impinge on the ability to meet a specific goal if not directly addressed.

Let's use a simple example to clarify

Imagine that a client identifies their goal as an income in retirement, and after careful consideration you establish that the income required is Rx per month after tax. If this client has not provided adequately for medical expenses in retirement (hopefully with medical insurance cover or medical aid), the income may be wholly inadequate if the client is faced with a major medical procedure. Financial advice can't consider meeting goals with appropriate investments in isolation, but must cover all other aspects of financial needs, like appropriate insurance cover.

It is also important to note that in most cases, we will be dealing with constraints that make meeting financial goals less than certain. It would be wonderful if assets existed to meet all conceivable goals, but the reality is far from this. The best we can hope for is for some assets that meet some of the dimensions of goals (liabilities) that we want to achieve, but even this can be rare and often it is also very expensive. Let's consider another simple example.

Imagine that an individual would like to save to buy a retirement home in the south of Portugal (the majestic Algarve) in 20 years time when they are planning on retiring. Portuguese inflation-linked bonds may seem like a good bet, but there are a couple of problems.

Firstly, the Portuguese government does not issue inflation-linked bonds. Secondly, even if they did exist (hypothetically), they may not have a maturity of 20 years, which would introduce either reinvestment risk (if maturity was less than 20 years) when they matured (and which will nevertheless exist on any coupon payments even if the maturity was 20 years), or market risk if the bonds would need to be sold before maturity (if maturity was substantially longer than 20 years). For clarity, market risk is the risk that the yield to maturity of the bond changes throughout its life (because of many factors that could affect the yield at different durations). The returns on bonds is therefore only somewhat known if they are held to maturity without default.

Thirdly, the Portuguese government could default on the (hypothetical) bonds before expiry, and the European Central Bank (ECB) may not provide any security on this default, in which case some of the capital invested would be lost (hopefully not all of it, although for governments this is a real possibility as bonds are not secured by assets).

Finally, and this is very important, house prices in the Algarve may not increase at the same rate as Portuguese inflation overall (which like all other country inflation is actually made up of a basket of goods consumed by an average household). In fact, given the demand for the majestic coast, house price inflation in the Algarve could be much higher. Clearly what started out as a simple goal, is actually not that simple and in actual fact can be very complex.

To explain goal-based investing, I will begin by describing the objectives of goal-based investing,

and the theoretically optimal solutions used to achieve these objectives, before introducing various constraints that will lead into the final proposal. The reason for tackling the topic thus, is to demonstrate the complexities that exist, so that the final solutions are not misunderstood as being simplistic, but rather as practical. It does however also aid with the understanding of the complexity that exists and that will remain within the final solutions in resolving financial goals.

Understanding goals

It is generally well understood that **goals can differ substantially along many different and important dimensions**. Let's begin by exploring these to understand the complexity that this creates. A goal can be of a capital nature (a single payment at some future date), or of an income nature (a couple of, or even many, payments at various future dates). They could be in one of many possible currencies (e.g. rand, dollar, or euro). They could be nominal amounts i.e. fixed amounts whose value is known today, or real amounts i.e. an amount linked to inflation (not only consumer price inflation, but any possible measure of inflation, say medical inflation). The future date/s could be in the near future (say next year), or very far into the future (say 80 years hence for a 20 year old starting their career and wishing to save for retirement and an income for when they are a centennial). The amounts required (capital or income) could be gross or net of tax. Finally, any of the above dimensions could be known with certainty, or completely uncertain e.g. how long will I live and need an income for? Do I know what education inflation will be and how it will relate to consumer price inflation? Do I know what future tax rates will be? Do I know what the price of a life annuity (with an insurance company providing longevity risk cover) will be? And so the list goes on.

It is important to distinguish between the dimensions of the goal(s) listed above, and the **investor preferences** associated with the goals, and priorities. Different investors may give different goals different priorities,

and it is important to understand and respect these when considering how to construct solutions to meet them. When certainty of meeting the specified goal is given the highest weighting (highest priority), the "best" solution will focus on minimising the risk of not meeting the goal. This solution is unlikely to be the same solution as when another dimension is given a higher priority e.g. maximising wealth or returns.

It is however appropriate to understand that the starting point to investing to meet a goal, is to find assets that best match the nature of the goals, as any deviation away from this introduces variability (uncertainty) in the outcomes. It is important to highlight that even the "optimal" solution would not be "risk-free" as uncertainty of meeting a specific goal could never be guaranteed (except in the simplest of cases). We'll touch on this in more detail below in the section on funding goals, but it should already be evident from the examples given in the introduction above.

The above complexity and conflicting objectives pose a serious problem which is not trivial to solve. To meet the various goals of various investors, you would need a large number of solutions to meet all of the different dimensions and requirements. While enough solutions probably exist globally to get you to a good answer for each goal, you would potentially need to know about and monitor hundreds or even thousands of these solutions. This is clearly untenable, which is why we begin by **simplifying the dimensions** so that we end up with a **manageable set of solutions** that will approximately match the dimensions of most goals.

Funding goals

Given the above as a starting point, how can we proceed? There are generally two approaches.

The first approach, which relates back to the theoretically "optimal" solution, would be to model the investments (every asset class available for investing i.e. not just theoretically available) and the goals (liabilities) stochastically i.e. randomly

(introducing their “known” uncertainty). We won’t go into the detail here around the modelling methodology, but it is important to know that modelling requires assumptions which are informed by historical data (or at least should be, and hence why I refer to it as “known”).

This means that the modelling considers all of the uncertainty around returns and correlations between asset classes, as well as the uncertainty of the liabilities and how they may change or evolve (and the correlation between the assets and the liabilities). The optimal solution to the problem is the investment (combinations of asset classes) that provides the best match to the liabilities (goal) i.e. the solution that minimises the uncertainty around meeting the goal, regardless of the cost. Investor preferences can then be factored into alternative solutions that deviate away from this based on other priorities.

The second approach, focuses on the traditional requirement for a discount rate to be used to equate the present value of the investments and goals (assets and liabilities), and does this deterministically i.e. not randomly. It is however important to note that even here stochastic modelling is used for the assets, but this is used to establish the expected risk of the assets and how to combine these efficiently (using a mean variance framework with expected return assumptions).

Mean-variance optimisation

Practically, this means building an “efficient frontier” i.e. a combination of assets that minimise variability (or variance/standard deviation) of returns for a given level of expected return. We actually do this in real return “space” (an abstract construct that looks at returns and variability in real terms i.e. after adjusting for inflation). This is straightforward when working in rand (because we can use SA CPI for inflation), but creates additional complexity when considering global goals and investments (which currency and inflation rate should be used?).

It should be obvious from the second approach, that you need a discount rate to equate assets (investments) and liabilities (goals). This creates the need for target returns in the solutions i.e. there is no alternative way (except for the first approach) to translate the goals into specific investment objectives.

Up to this point, things should be fairly clear. We’ve constructed an efficient frontier that provides us with the combination of assets to be used for a given level of risk or return (real). It is important to appreciate the sensitivity of the results obtained, to the assumptions made (and methodology adopted), so that we don’t get too comfortable with the “preciseness” of the numbers i.e. it would be wrong to think of this efficient frontier as being “certain” in any meaningful way as one of the dimensions actually represents uncertainty (risk).

The portfolios on the efficient frontier will then have a corresponding asset allocation to the asset classes used in the modelling, and this is used as the strategic asset allocation (SAA) for the corresponding solution. All that is left to do, is decide on how many solutions are needed, and exactly where along the frontier should these be selected from, and we will consider this next.

Mean-variance optimisation may appear dated, but it remains a useful and powerful tool in understanding how to build portfolios under certain assumptions. It can incorporate Monte Carlo simulations using historical data, or parametric distributions based on historical data. It can incorporate historical or expected returns, and can even incorporate stochastic covariances (correlations) between the various asset classes i.e. uncertainty can be introduced into the various dimensions of interest.

Building solutions

So how do we move from the optimisation work, and the resultant possible solutions, to a range of portfolios to meet a varied range of individual goals? One obvious extreme method would be to include just one portfolio (somewhere on the frontier), and force everyone to use this portfolio for every goal.

Clearly this is not very client-centric, and appears to be a little too extreme in terms of simplification. Another less obvious extreme may be to have many portfolios (say 50) all along the frontier, hoping to provide a lot of granularity in meeting various risk and return requirements. We hope that it is obvious that this is not practical or necessary, and actually highlights a lack of understanding of the uncertainty present in modelling and dependence on the assumptions.

So, having thrown out the extremes, we can focus on finding a suitable compromise, but let's consider the compromise I'm discussing in more detail first. Too many portfolios are extremely costly to manage (on various cost dimensions, including indirect costs related to governance), and we therefore want to minimise the number of portfolios to minimise these costs, costs which will need to be passed on to clients. Too few portfolios on the other hand, don't provide enough granularity in terms of meeting different risk and return requirements. This is what we will need to balance, and find a reasonable compromise around.

On the lower limit, we could build just two portfolios (100% local cash, and say 100% local equities), and every client could be given a combination of these two to meet their specific requirements. It is important to understand the limitations of this possible solution. The first, is that it is sub-optimal in a mean-variance sense (i.e. the combination will not lie on the efficient frontier except for the two extreme cases) because it doesn't make use of all available asset classes, which provide diversification benefits. The second, is that it could be sub-optimal from a tax and cost perspective because it would require constant rebalancing to maintain a fairly constant allocation to cash and equities.

If we consider 1% real return increments from 1% for local cash to 7% for equities (approximately our long-term real expected return assumptions), we could end up with five multi-asset class portfolios ranging from 2% to 6%, giving us a good range of

portfolios to meet most investors' risk and return requirements (in addition to cash and equities at the two extremes for investors looking for something more). Some people may argue for even greater granularity (i.e. more portfolios at say 0.5% increments), but the above proposal already introduces spurious accuracy i.e. there is already so much uncertainty around what each portfolio will deliver over various time frames.

It is important to understand that there is no "correct" or "optimal" number of portfolios, or where they should be positioned on the efficient frontier. If we are given a specific utility function that captures the client's preference with reference to competing constraints, it is fairly simple to point to an optimal solution, but generally this is derived through a conversation with clients around the priority of the competing objectives and constraints. To suggest otherwise demonstrates a lack of understanding, and is simply misleading.

Mapping goals to solutions and understanding the limitations

Now that you have a range of portfolios along the dimensions of expected risk and return, you need to decide on which portfolio to use to meet each specific goal. I will deliberately sidestep the issue of whether investments should be considered separately for each individual goal (as opposed to collectively which is actually more optimal) as this remains a contentious issue and difficult for many to grasp.

This is where the traditional approach of "risk-profiling" investors enters the advice framework, although I think this will ultimately evolve away from this (another contentious issue I will avoid in this article). The traditional approach considers three dimensions of risk, which includes risk capacity, risk required and risk tolerance (the dimension where psychological questionnaires are used to establish attitude to and appetite for risk).

It is critical for the investor to understand risk as uncertainty at this point, as many investors may believe that this approach to investing for meeting goals removes all uncertainty, where nothing could be further from the truth. The methodology actually enables a discussion around the dimensions of goals and investments, and their attendant uncertainties, so that an appreciation of the complexity can be reached. Financial advisers will be doing their clients a great disservice if they don't use the opportunity to have this discussion upfront as they may learn later when their clients become disgruntled by "poor" performance.

This is where the use of great tools/aids can assist financial advisers and their clients in understanding these dimensions and risks, and graphical representations of the evolution of the investment and the goal can be very enlightening. Scenario analysis and hypotheticals are two more great tools to help in this mammoth task e.g. showing how the investment would have performed through the global financial crisis (GFC). If a client is uncomfortable with the level of drawdown through the GFC, they should seriously consider lower risk portfolios as this scenario could easily repeat in the lifetime of the goal.

A tool that allows an adviser to flex (change) various dimensions associated with the investor, their goal, and possible solutions, is extremely powerful in trying to find a suitable investment to meet an investor's very specific requirements. The investor should be able to see (visually as well as understand) the impact of changing the investment or consumption horizon, the initial and ongoing investment contributions (if any), the expected risk and return assumptions, and the certainty (probability of achieving the goal), on the goal value. The investor should then be able to change the question around to ask what the impact would be on any of those same dimensions, if the goal value were changed e.g. if the investor wants a higher amount at retirement, how much longer should the investor work before retiring?

Evaluating performance and ongoing investment advice

Once all of the above has been adequately covered, with the investor demonstrating a good understanding of the methodology and how it will assist in meeting their specific goals, a record of advice can be produced for both the investor and the financial adviser. It is critical that this record includes the uncertainty discussed as this is one of the most important dimensions of the exercise and will represent the most discussed issue in the annual review of how the investment is tracking against the goal. It would be simple if the trajectory of the investment progressed smoothly along the expected return path, but this is not only unlikely, but actually practically impossible.

At each review, the financial adviser can therefore consider how far above or below the trajectory the investment is progressing, and whether any corrective action should be taken. There are many things to consider in this process, so I will not be tackling them here, but it again represents a wonderful opportunity for adviser and investor to have a discussion around the initial process and their shared understanding of how the investment would evolve. By spending adequate time doing this, it should prevent any short-term irrational decisions that could be detrimental to the long-term success of meeting goals, which was the initial intention of following this methodology.

Conclusion

It is important to recognise what goal-based investing aims to achieve, and the idealised solutions that would theoretically be employed to meet them. It is equally important to understand the practical considerations that are needed when arriving at real world solutions, and the limitations and compromises that have been made to arrive at these. It is then fairly easy to understand why the solutions look the way they do, and how that can be integrated into an advanced financial advice

framework. Without this understanding, it is easy to criticise the solutions as simplistic, and advisers should be aware of this, so that they can defend the methodology and approach to their clients.

I have taken care to articulate these complexities and discuss appropriate ways of addressing them before presenting a solid foundation for the methodology and recommendations made. I would urge all stakeholders to put sufficient emphasis on this understanding, before embarking on this methodology of investing to meet goals. I think that clients want to see consistent and integrated thinking and advice, and goal-based investing is well positioned to provide it, but requires a deep understanding of the complexity and the time to get the client to a good level of understanding.

The time invested upfront will be worth it as the adviser meets with clients annually along the journey, comparing how the investments are tracking relative to the goals. This presents a wonderful opportunity to stop the “short-termism” prevalent in the industry as investors chase the best past performers according to some survey or peer group ranking tables, in the belief that past performance may in fact be a good guide to future performance, despite all the “health” hazards communicated around this.

It is important to recognise what goal-based investing aims to achieve, and the idealised solutions that would theoretically be employed to meet them. It is equally important to understand the practical considerations that are needed when arriving at real world solutions, and the limitations and compromises that have been made to arrive at these.