Lecture Note 3.3

Goal-Based Investing

Individual investors' problems can be broadly summarized as a combination of various wealth and/or consumption goals that an investor wants or needs to achieve, subject to strict dollar budget constraints, defined in terms of initial wealth and future income, as well as risk budget constraints such as maximum drawdown limits for example. At the starting point of the goal-based wealth management approach lies the recognition that investors need *solutions* to these problems, as opposed to off-the-shelf investment *products*, which are mainly marketed on the basis of their performance. It is indeed important to note that the success or failure in satisfying these goals subject to dollar and risk budgets does not critically depend upon the stand-alone performance of a particular fund nor that of a given asset class. It depends instead upon how well the performance on the investor's portfolio dynamically interacts with the risk factors impacting the present value of the investor's specific goals. In this context, one may argue that the key challenge for financial advisors is to implement dedicated investment solutions aiming to generate the highest possible probability of achieving investors' goals, and a reasonably low expected shortfall in case adverse market conditions make the achievement of said goals unfeasible. Taking investors' goals as inputs to design a dedicated investment solution is known as *goal-based investing* (GBI in short), which is the exact counterpart to *liability-driven investing* in institutional money management, where investors' problems are summarized in terms of their liabilities, as opposed to goals.

More precisely, the focus of a meaningful goal-based investing solution is to initiate an investor-centric approach to the investment management process which should lead to maximizing the probability of investors reaching their *target* levels of wealth/consumption, also known as *aspirational goals*, while securing their *minimum* levels of wealth/consumption, also known as *essential goals*. Achieving a high probability to reach the objective always requires the combined use of three main ingredients: (i) a set of safe *goal-hedging portfolios* dedicated to the protection of the essential goals; (ii) a well-diversified risky *performance-seeking portfolio* that provides upside potential by efficiently harvesting risk premia across and within asset classes; (iii) a dynamic allocation to the safe and risky building blocks that allows investors to reliably protect essential goals in poor market conditions while generating the highest probabilities of reaching the aspirational goals in more positive market environments.

Each of these items corresponds to one of the risk management principles that have been introduced in the previous Lecture Notes: (i) hedging (of the risk factors relevant to the essential goals); (ii) diversification (of unrewarded and rewarded risk exposures); and (iii) insurance (to protect against downside risk relative to the essential goals while keeping the opportunity for upside potential needed to reach the aspirational goals). These sources of value are already used to some extent in different practical contexts, but a *comprehensive* disciplined investment management framework is required for the design of investment solutions adapted to investors' problems. This Lecture Note is precisely intended to provide to a step-by-step description of such a goal-based investing framework. For this, we heavily rely on the material introduced from the previous Lecture Notes with relatively little emphasis of the technical content in order to focus instead on applications. An in-depth description of the theory and practice of goal-based investing is given in Chhabra (2005) and Deguest et al. (2015).

While ultra high net worth individual clients can possibly benefit from a dedicated fully customized goal-based investing solution, it should be emphasized that the framework aims to extend well beyond these situations and can also be applied in the context of suitably designed mass-customized investment solutions to relatively homogenous clusters of affluent, mass affluent or retail clients who share some common types of goals. The most important goals for most individuals or households are undoubtedly financing college education, financing home acquisition, and financing consumption needs in retirement. In this Lecture Note, we present an example of mass-customized goal-based investing principles applied to the retirement context and refer to Martellini and Milhau (2016) for a more detailed treatment.

3.3.1 Step 1: *Identification of Investors' Goals*

An adapted goal-based investing process must always start with a proper identification of an investor's hierarchical set of goals, as well as available resources which can be best summarized in terms of the investor's "balance sheet". It is important to emphasize that the collection of such detailed pieces of information on an investor's situation are the key inputs in a goal-based investing process. In current wealth management practice, however, such information is either not collected or collected but not used in any relevant way. Even for high net-worth individuals, who can physically engage in an in-depth dialogue with a personal advisor, there is little evidence that the wealth of information collected in this process is directly used in the design of dedicated investment solutions. The rise of robo-advisors opens new perspectives by giving the opportunity to reach out to a large number of individuals, regardless of their wealth levels, and to collect useful information through digital platforms in a highly scalable and cost-efficient manner, without the need for human intervention.

Formally, a goal can be described by a set of dates and a set of values representing the desired levels of wealth or payment. In the important application to retirement saving, an investor's goal is to generate replacement income in retirement and the cash-flow dates are determined by the worker's age and projected retirement date. In a general context, many different goals may co-exist at any point in time, in which case they need to come with a priority ranking. In practice, however, it is easier to identify a given goal for one investor, of several similar investors, and design a dedicated investment solution to accommodate this particular goal.

3.3.2 Step 2: *Classification of Investors' Goals*

Once a list of goals has been obtained from the client, the next step is to establish a hierarchical classification of these goals in two main categories, affordable and non-affordable. [1] A discussion of affordability conditions has been given in Lecture Note 2.1, and we simply recall here that affordable goals are goals

[1] In the presence of multiple goals, this leads to the definition of the *maximal set of affordable goals*. It is the largest set of goals that are jointly affordable, constructed by starting from the most important one, adding the other ones by order of decreasing priority and stopping as soon as the required capital exceeds the available resources.

that can be secured with a 100% probability with some suitable strategy given current wealth and future income. Another important distinction exists, related to the affordable versus non-affordable classification, which is expressed in terms of the classification of goals in essential versus aspirational goal. To understand this additional distinction, let us first recognize that even if a goal is affordable, an investor may decide not to secure it with probability 1 in case the goal is not perceived as important enough to justify committing the resources required to securing it. An *essential goal* is precisely defined as an affordable goal that the investor wants to protect with 100% probability. It can also be regarded as a *floor* for the associated allocation strategy as defined in Lecture Note 3.2 in the sense that it represents a minimum level of wealth that the strategy should guarantee in all scenarios. Conversely, if a goal is perceived as vital to an investor but it is not affordable, then the investor will have to regard it as an aspirational goal, which can potentially be achieved but not with probability 1, and the investor can set a lower level for the same goal as an essential goal. More generally, non-essential goals are said to be *aspirational*, and they are of two kinds. They can be goals that would be affordable (together with other essential goals) but that the investor decides not to secure, or they can be non-affordable goals. A typical example is that of an individual in the accumulation phase who sets both a minimum level and a target level of replacement income in retirement. If the minimum level is affordable, then it can be treated as an essential goal, while the target level will be treated as an aspirational goal assuming it is not affordable.

Now, why would investors give up on protecting some goals that are affordable? In fact, if all goals are jointly affordable, there is no reason why they would. As long as they have enough resources available (under the form of current wealth or future income) to secure all of their identified goals, then they should rationally secure all of them. Unfortunately, this is of course almost never the case since even ultra high net worth individuals also have non-affordable goals, including charitable goals (think for example about Bill & Melinda Gates foundation and their effort to save million of lives in Africa, a non-affordable goal even for the world wealthiest individual). It is critically important to recognize that a trade-off always exists between security and upside potential, and this because securing essential goals involves an opportunity cost in terms of upside potential. On the one hand, individuals indeed need to invest a substantial fraction of their assets in a performance-seeking portfolio (PSP) to reach ambitious aspirational goals. On the other hand, securing essential goals involves setting floors, which eventually reduces risk taking in asset allocation decisions. Thus, increasing an essential goal level will lead to a higher protection level, but it implies a lower allocation to the PSP, and as a result a lower probability of reaching aspirational goals.

Figure 3.3.(a) gives a representation of this trade-off between security and upside potential in the context of a simple CPPI strategy. The CPPI strategy aims to protect a fixed level of wealth at the ten-year horizon. The essential goal is expressed as a percentage δ_{ess} (necessarily less than 100%) of the maximum wealth level that is affordable at date 0: this means that the minimum wealth to protect at the horizon is

$$\delta_{ess} \times \frac{W_0}{b_{0,10}},$$

where W_0 is the initial wealth and $b_{0,10}$ is the initial price of the ten-year zero-coupon bond. Similarly, the aspirational goal is a percentage greater than 100% of this maximum. For low levels of the essential goal, probabilities of reaching the aspirational goal are relatively insensitive to the specification of this goal, but they are decreasing for higher levels, and they eventually fall to zero when the investor wants to protect

Figure 3.3.(a): Probability of reaching an aspirational goal as a function of the essential goal.

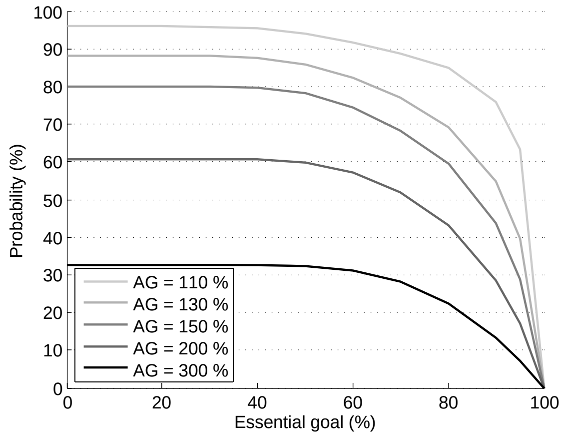

Note: PSP volatility = 21.38%; PSP Sharpe ratio = 0.40; short-term rate = 3.06%; correlation between short-term rate and PSP = 0. The success probability is the probability of reaching the aspirational goal at least once within the investment period. The portfolio is rebalanced quarterly with a multiplier of 3.

100% of the affordable maximum.

It is clear from the figure that in order to increase the probabilities of reaching high goals, one has to set the floor to a lower level. But this comes at the cost of lower secured wealth levels. For instance, in order to reach 130% with a minimum 75% probability, one has to set the floor to about 70%. In other words, even if he is in principle able to choose a protection level of 100%, an investor may decide to choose a lower floor in order to have a decent probability of arriving some day at the non-affordable level of 130%. Those who are not willing to underperform the ten-year zero-coupon bond by more than 30% must either consider a lower aspirational goal (e.g., 110%) or satisfy themselves with a lower success probability. While this example is highly stylized, it conveys a very general and important message. By setting higher levels for their essential goals, investors may improve outcomes in the worst scenarios, but they also reduce their ability to reach non-affordable goals. This is the reason why any meaningful dialogue with an investor must be based on a tool allowing the investor to analyze and measure the opportunity costs associated with their set levels of essential goals. In the absence of this analysis, investors might be tempted to set for themselves unreasonably high levels of protection, with an exceedingly large associated opportunity cost in terms of upside potential.

3.3.3 Step 3: *Identification of Investable Building Blocks*

The third step in the framework is the design of building blocks dedicated to the achievement of goals. As already mentioned in the introduction, the fund separation theorem presented in Lecture Note 2.2 (Proposition 2.2.A) provides formal grounding to the intuition that suggests a separation between a "safe" goal-hedging building block, the role of which is to secure essential goals, and a risky performance-seeking building block, which is expected to deliver the kind of upside performance needed to achieve aspirational goals.

Goal-Hedging Portfolios

To each essential goal is associated a suitably designed goal-hedging portfolio (GHP), which is the exact counterpart in individual money management to the liability-hedging portfolio in institutional money management (see Lecture Note 2.2). In contrast to PSPs, GHPs are not well-diversified portfolio of rewarded factor exposures; they instead tend to be concentrated portfolios loading on the same risk factors as the present values of the goals. In the case of the retirement goal, where the individual seeks to secure minimum (essential) and target (aspirational) levels of consumption in retirement, the GHP is an inflation-linked deferred annuity that starts paying inflation-linked replacement income on the retirement date. Now, if this deferred annuity is the perfect GHP for investors preparing for retirement, it turns out that in practice the demand for annuities is low, a fact referred to as the "annuity puzzle". Various explanations have been put forward (see Pashchenko (2013) for a survey). In particular, the irreversibility, cost inefficiency and opacity of annuities are perceived as a major obstacle. In view of these problems, individuals saving for retirement would benefit from holding a bond ladder as a *retirement goal hedging portfolio*, which would deliver replacement income for a fixed period in retirement (say the first 20 years of retirement), and eventually purchasing late life annuities for taking care of their consumption needs beyond that period.

Performance-Seeking Portfolio

As already mentioned, performance-seeking portfolios (PSPs) are well-diversified portfolios allowing for the most efficient harvesting of risk premia across and within asset classes. Diversification (as opposed to hedging) is indeed the risk management technique that allows investors to eliminate or at least reduce (diversify away) unrewarded risk in their portfolios, which allows them to enjoy higher rewards per unit of risk, and therefore a higher average funding ratio at horizon for a given risk budget.

In theory, the proper diversification criterion to maximize is the Sharpe ratio, but it has long been recognized (see Lecture Note 1.1) that expected returns are extremely difficult to estimate, so that Sharpe ratio maximization is usually replaced by alternative approaches involving some mixture of naive and scientific diversification (see Lecture Notes 1.2 and 1.3). In particular, the *factor investing* and *smart beta* paradigms in the equity universe have led to the successful development of *smart factor indices*, whose objective is to give full access to the premium of a given rewarded factor, such as value, size, momentum, volatility, investment or profitability, among others.

3.3.4 Step 4: *Identification of an Allocation Strategy to the Building Blocks*

Once the building blocks of the solution have been identified, the next question is to design, calibrate and implement an allocation strategy that secures the essential goal while giving high probabilities of reaching the aspirational ones.

Throughout this section, we focus on the simple case where there is only one essential goal translating into a minimum wealth constraint of the form $W_t \geq F_t$, where F_t is the floor. To fix the ideas, the investor could be an individual saving for retirement, in which case the essential goal is to secure a replacement income equal to a fraction δ_{ess} (necessarily less than 1) of the maximum income that is affordable at the beginning of the investment period (date 0). Then, the floor is

$$F_t = \delta_{ess} \frac{W_0}{\beta_0} \beta_t,$$

where β_t is the price of a deferred annuity with cash flows normalized to \$1. If we decide to assume away longevity risk, or only secure replacement income for a fixed period of time in retirement, the GHP is a portfolio of pure discount bonds with maturity dates matching the dates when replacement income is needed.

Maximizing Success Probability

Suppose that the investor has a well-defined aspirational goal, expressed as a multiple α greater than 1 of the floor. In the retirement example, the aspirational goal is a level of replacement income equal to a multiple δ_{asp} of the initially affordable income, so the present value of the aspirational goal is

$$G_t = \alpha F_t, \tag{3.3.1}$$

with $\alpha = \delta_{asp}/\delta_{ess}$.

A natural objective is to maximize the probability of reaching the aspirational goal by the horizon T, subject to the constraint of securing the essential one with probability 1. This problem is solved for example by Browne (1999) and Martellini and Milhau (2016) (see also Föllmer and Leukert (1999)). In a Black-Scholes framework with deterministic interest rates, volatilities, risk premia and correlations across assets, the optimal strategy involves a dynamic allocation to the MSR performance-seeking portfolio, the GHP and the cash account. Let us assume that the goal is proportional to the cap, like in Equation (3.3.1). Then, the optimal weights of the risky assets are:

$$\boldsymbol{w}_t = \varphi_t \frac{\lambda_{MSR,t}}{\sigma_{MSR,t}} \boldsymbol{w}_{MSR,t} + (1-\varphi_t)\, \boldsymbol{w}_{GHP,t},$$
$$\varphi_t = \frac{(\alpha-1)F_t}{\sigma_t W_t} n\left[\mathcal{N}^{-1}\left(\frac{W_t - F_t}{(\alpha-1)F_t}\right)\right],$$

where n is the Gaussian probability distribution function, $\mathcal{N}$ is the Gaussian cumulative distribution function, σ_t is the standard deviation of $\log\left[W_{go,T}/F_T\right]$ evaluated at date t, and $W_{go,T}$ is the payoff of the *growth-optimal strategy*, which is the optimal unconstrained strategy of Lecture Note 2.3 with unit risk-aversion.

The payoff associated with the optimal strategy can be regarded as a the result of combining the floor with a binary option, since it is given by:

$$W_T = F_T + (G_T - F_T)X,$$

where the indicator function X is 0 or 1: it is 1 if the *growth-optimal strategy* outperforms the quantity $G_T - F_T$ by a sufficient amount. In details, the success condition – which makes X equal to 1 and W_T equal to G_T – is:

$$\chi W_T^{go} \geq G_T - F_T,$$

Figure 3.3.(b): Allocation to MSR portfolio in capped OBPI and probability-maximizing strategies.

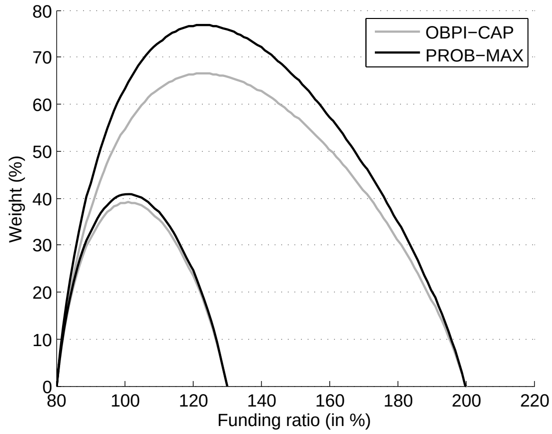

Note: Horizon = 10 years; volatility of growth-optimal strategy = 15% per year; initial floor-to-wealth ratio $F_0/W_0 = 80\%$; initial cap-to-wealth ratio $G_0/W_0 = 130\%$ or 200%. It is assumed that $\lambda_{MSR,t} = \sigma_{MSR,t}$.

where χ is a threshold greater than $1 - F_0/W_0$. The coefficient χ must be computed numerically by writing that initial wealth equals the floor plus the binary option price. In the Black-Scholes setting, we have an explicit expression for χ:

$$\chi = \left[\delta_{asp} - \delta_{ess}\right] \exp\left[\sigma_0 \mathcal{N}^{-1}\left(\frac{1-\delta_{ess}}{\delta_{asp}-\delta_{ess}}\right) + \frac{\sigma_0^2}{2}\right],$$

where σ_0 is the standard deviation of $\log\left[W_{go,T}/F_T\right]$ and $\mathcal{N}^{-1}$ is the inverse of the Gaussian cumulative distribution function.

The probability-maximizing payoff is somewhat similar to the payoff of the capped OBPI approach, in that they both are bounded from below by the floor (essential goal) and from above by the cap (aspirational goal). The key difference is that the probability of reaching a wealth level between the floor and the cap is strictly positive for the capped OBPI approach while it is zero for the probability-maximizing strategy. Indeed with the latter strategy any value strictly lower than the cap is a failure to reach the goal, and should therefore be avoided as often as possibles.

Figure 3.3.(b) shows the allocations to the MSR portfolio in the probability-maximizing strategy and the capped OBPI with a risk aversion parameter equal to 1 (see Section 3.2.5 in Lecture Note 3.2). For a higher risk aversion level γ, the former strategy would be unaffected by definition, and the weight allocated to the MSR building block in the latter would decrease. On the x-axis is the funding ratio, defined as:[2]

$$R_t = \frac{W_t F_0}{F_t W_0}. \tag{3.3.2}$$

The conclusion from the figure is that both types strategies embed a "smooth landing" mechanism on the floor and the cap, with an allocation to the MSR portfolio that gradually shrinks to zero as wealth approaches

[2]Note the difference with respect to the wealth-to-floor ratio, equal to W_t/F_t.

the bounds.

The growth-optimal portfolio can be replaced with any other risky asset. This implies a different coefficient χ, and, in general, a lower success probability. It is only when the growth-optimal strategy is used as the underlying asset that the probability is maximum.

Implementable Strategies

The probability maximization objective provides useful guidelines for the design of goal-based investing strategies, but the theoretically optimal strategy has shortcomings that make it unattractive or even unavailable for practical applications. First, it typically involves unreasonably large levels of leverage. Second, the dynamic replication strategy for the optimal payoff exhibits an explicit dependency on the Sharpe ratio of the MSR portfolio, which is not an easy to estimate parameter.

We now seek to identify a simple rule-based strategy that should enjoy some of the desirable properties of the probability-maximizing policy, including in particular secure the essential goal with probability 1 and generate a substantial probability of reaching the aspirational goal, without the associated drawbacks. In particular, we would require the strategy to satisfy no leverage constraints, and also be based only on observable quantities, or at least on easily estimable parameters. We can use for this an investment rule with a CPPI flavor, where the allocation to the risky building block is a function of the distance between the current wealth and the current floor value (both of which are observable), and also on the multiplier, which is a technical parameter specified by the portfolio manager. In details, the fraction of wealth allocated to the performance-seeking portfolio is:

$$\pi_t = m\left(1 - \frac{F_t}{W_t}\right),$$

with the usual truncation to 0 and 100%. The remainder goes to the GHP. A useful refinement in implementation consists in taking a time-varying multiplier, so as to recognize that investors may benefit from increasing (respectively, decreasing) the aggressiveness of their strategy in periods where the PSP has historically low volatility and high risk premia (respectively, historically high volatility and low risk premia).

We may call these extended CPPI strategies *goal-based investing* (GBI) strategies since they are designed to secure essential goals. It should be noted that the investment rule is independent from any aspirational goal level. To achieve a higher probability of reaching such a goal, it is desirable to implement an additional *stop-gain* mechanism by which the assets are transferred to the corresponding GHP at the first hitting time by the wealth process of the present value of the goal.

Mass Customization: The Retirement Example

We now briefly discuss how the strategy can be made scalable in the retirement saving context. Generating replacement income in the decumulation phase of their life is a concern for every person in accumulation, so this retirement investing is the archetypical situation where mass-customization can be useful.

To fix the ideas, consider an individual who brings an initial contribution of $10,000 at date 0 and plans to retire in 20 years, and assume that the price of the deferred annuity that pays $1 every year in retirement (or a bond ladder that pays $1 every year for the first 20 years in retirement) is $10. The maximum affordable income that the investor can secure given the initial wealth level is thus $1,000 per year. The problem for this single individual is relatively simple, and can be summarized as the need to protect a fraction – say 80% – of this annual income, and to have a high probability – say at least 75% – of increasing it to a higher aspirational level – say $1,300. This problem can be addressed with a GBI strategy with a floor equal to the price of a deferred annuity (or bond ladder) that pays $800 per year. This strategy defines a retirement investment fund characterized by an inception date (the beginning of accumulation for this individual) and a terminal date (the retirement date).

Now, consider other individuals with the same date, but arriving at some arbitrary dates later down the road. To provide all individuals with the same guarantee – that is with the same 80% fraction of their initially affordable income –, the fund must outperform the goal-hedging portfolio (deferred annuity bond ladder) by at least 80% between *any* date and the terminal date. This is mathematically equivalent to imposing that the fund value, W, must satisfy:

$$W_T \geq 0.8 \times \beta_T \times \max_{t \leq T} \frac{W_t}{\beta_t}.$$

The right-hand side of this equation defines a "relative maximum drawdown" floor: the fund value expressed in the annuity numeraire must never fall below 80% of the maximum ever attained.

As a conclusion, it is possible to accommodate individuals with different contribution levels and different arrival dates with a single fund, in which the floor is of the relative maximum drawdown type. Introducing a relative maximum drawdown floor, however, implies that the floor never decreases relative to the present value of the goal; it can only increase or stays flat relative to the goal value. As a result of this increase in the floor level, available risk budgets are lower from the perspective of initial investors than they would have been with a standard floor taken as a constant fraction of the goal, which induces an opportunity cost that can and should be measured. More generally, one may also envision launching new funds if and when risk taking has become too small for existing funds to generate a sufficient probability for new investors to achieve reasonably high aspirational goal levels (see Martellini and Milhau (2016) for details).

3.3.5 Step 5: *Introduction of a Goal-Based Reporting Process*

The goal-based investing framework should also encompass a process dedicated to facilitating an efficient dialogue with the investor. For a goal-based investing solution, standard risk and return indicators such as historical performances, volatilities and Sharpe ratios are indeed mostly irrelevant, and they should be replaced by more meaningful information on the goals that have already been achieved and on the probabilities of reaching those that have not yet been reached.

Achieved Goals and Protected Levels

To each GBI solution is attached an essential goal that translates into a floor. Up to gap risk – which should remain small in likelihood and magnitude if the multiplier has been properly specified –, the fund value stays above this floor at all times. The minimum level of wealth guaranteed by the strategy at the horizon can thus be converted into an indicator that can be expressed in a meaningful way with respect to the investor's goal. If the goal is to protect a fixed level of wealth, the indicator to be reported is simply this secured wealth level. In case of retirement investing, the relevant information is the level of replacement income that can be secured given the initial wealth (as well as future contributions, if any). Here, the relative maximum drawdown floor, which has been introduced for mass customization motives, proves to be useful from the perspective of a single investor. Indeed, we have, for any date t:

$$\frac{W_T}{\beta_T} \geq \delta_{ess} \times \max_{s \leq t} \frac{W_s}{\beta_s},$$

where δ_{ess} is 80% for instance. This means that the replacement income on which the investor can rely for date T is at least 80% of the maximum replacement income obtained until date t. The protected level to report is 80% of the maximum, an indicator that never decreases and grows both because of the fund performance and also because of the contributions made by the individual down the road. This indicator has the advantage of being completely observable at the reporting date, and it does not require any assumption on the future performance of the fund building blocks, nor on the future contribution scheme.

Together with the protected level should also be reported the value of the achieved goal. In the retirement example, this is W_t/β_t, the currently affordable income. The current situation can be summarized in the funding ratio (3.3.2), which by definition is 100% at the initial date and should always be greater than F_0/W_0. The information on the funding level can be used to trigger the stop-gain rule as soon as the aspirational goal is attained.

Success Probabilities

As long as aspirational goals are not attained, a natural question is how likely they are to be reached within the remaining time. This question can typically be addressed through the use of Monte-Carlo simulations. The stochastic model should include a set of assumptions on the dynamics of the risk factors that affect the value of the fund and the present values of the aspirational goals, as well as parametric assumptions for the dynamics of these risk factors. At this stage, model risk and parameter risk are necessarily involved, but it is important to emphasize that they only impact the reported probability of achieving aspirational goals, and not the strategy itself, which is only based on observable quantities.

Beyond the dynamics of the financial state variables, one can also make assumptions on the future contributions made by the investor. This is particularly relevant in the retirement saving context, where accumulation can start early in the life cycle and the initial capital is low compared to the cumulative value of expected future contributions. It can be useful to report both a probability based on the assumption of zero future contributions and one based on a positive constant contribution level. Measuring the positive impact of contributions on the probability of success is an extremely powerful way to provide incentives for investors to start preparing for retirement sufficiently early in their life cycle, well before the transition stage.

References

Browne, S. (1999). Reaching goals by a deadline: Digital options and continuous-time active portfolio management. *Advances in Applied Probability 31*(2), 551–577.

Chhabra, A. (2005). Beyond markowitz: A comprehensive wealth allocation framework for individual investors. *Journal of Portfolio Management 7*(5), 8–34.

Deguest, R., L. Martellini, V. Milhau, A. Suri, and H. Wang (2015). Introducing a comprehensive risk allocation framework for goals-based wealth management. EDHEC-Risk Institute Publication.

Föllmer, H. and P. Leukert (1999). Quantile hedging. *Finance and Stochastics 3*(3), 251–273.

Martellini, L. and V. Milhau (2016). Mass Customisation versus Mass Production in Retirement Investment Management: Addressing a "Tough Engineering Problem". Working paper.

Pashchenko, S. (2013). Accounting for non-annuitization. *Journal of Public Economics 98*, 53–67.