UvA-DARE (Digital Academic Repository)

A multi-objective decision framework for lifecycle investment

Timmermans, S.H.J.T.; Schumacher, J.M.; Ponds, E.H.M.

DOI

[10.2139/ssrn.3038803](https://doi.org/10.2139/ssrn.3038803)

Publication date

2017

Document Version

Submitted manuscript

[Link to publication](#)

Citation for published version (APA):

Timmermans, S. H. J. T., Schumacher, J. M., & Ponds, E. H. M. (2017). *A multi-objective decision framework for lifecycle investment*. University of Amsterdam, section Quantitative Economics. <https://doi.org/10.2139/ssrn.3038803>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A multi-objective decision framework for lifecycle investment ^{*}

Sjoerd H.J.T. Timmermans[†] Johannes M. Schumacher[‡] Eduard H.M. Ponds[§]

September 18, 2017

In this paper we propose a multi-objective decision framework for lifecycle investment choice. Instead of optimizing individual strategies with respect to a single-valued objective, we suggest evaluation of *classes* of strategies in terms of the quality of the *tradeoffs* that they provide. The proposed framework takes inspiration from psychological theories which, on the one hand, assert that humans analyze risky choice situations in terms of several competing factors, and, on the other hand, recognize that attribute overload is detrimental to decision making. In particular, we use SP/A (security-potential/aspiration) theory as developed by Lopes and co-authors. The proposed approach is illustrated in a simple lifecycle model. As decision factors, we consider (a) the contribution paid, (b) the ambition level (targeted level of retirement income), and (c) the guarantee level (a level of retirement income that will be achieved with high probability). In terms of the tradeoffs generated between these indices, we compare a class of traditional lifecycle strategies, defined in terms of a glide path, with a class of so called collar strategies.

Keywords: pension plan design, choice architecture, lifecycle investment, dynamic asset allocation.

^{*}We thank Jan Snippe and Ralph Koijen for helpful comments on an earlier version of this paper.

[†]Email: sjoerdtimmermans@gmail.com

[‡]University of Amsterdam, section Quantitative Economics, email: J.M.Schumacher@uva.nl

[§]APG, Tilburg University, and Netspar, email: eduard.ponds@apg-am.nl

1 Introduction

It has been documented extensively that many people have difficulties in making good choices for retirement planning; see for instance Benartzi and Thaler (2007); Lusardi and Mitchell (2011); Mitchell and Utkus (2004)). Behavioral biases play a role even among trustees of pension funds (Clark et al., 2006). Considerable attention has been paid in the literature to ways in which such biases may be mitigated in the face of financial illiteracy and irrationality (cf. for instance Thaler and Benartzi (2004); Bodie and Prast (2012)). In this paper, we address a related but different question, namely how decisions on investing for retirement may be made by a person who is well informed and who is willing to spend time to arrive at a good decision. In other words, our aim in this paper is not so much to provide support for the weak, but rather to provide support for the strong. We have individuals in mind whose level of financial informedness and sophistication would be comparable to, say, that of a qualified trustee, but who would not necessarily be able to describe their preferences in terms of fully specified utility functions and probabilities across future events.

When the model of rationality is accepted that has been proposed by Savage (1954), it is in principle clear how to take a rational decision under uncertainty: form a (subjective) probability assessment, specify a utility function over outcomes, and then maximize expected utility. In the context of lifecycle investment decisions, this format has been followed in a broad stream of academic literature starting from the seminal work of Merton (1969, 1971). In the actual practice of decision making by pension fund boards or by committed individuals, the model has however not obtained great popularity. One reason for this may be that pension fund boards as well as individuals typically do not find it easy to specify a utility function. A more fundamental reason could be that individuals, even when well informed, do not always take their decisions in line with Savage's model of rationality, for instance due to the presence of ambiguity (Ellsberg, 1961).

Selection of an optimal probability distribution of retirement income is further complicated by the fact that the object of interest has many degrees of freedom, and may therefore present a case of choice overload. The demotivating effect of presenting too many options in pension plan decisions is well known (Iyengar et al., 2004), but one may argue that this effect should be overcome by a dedicated individual or by a board of trustees. Evidence from the psychological literature suggests however that the presence of many attributes can lead to choice situations that surpass the cognitive abilities of humans (Fasolo et al., 2007). Forms of decision support would therefore be called for. General methodologies for decision aid are discussed for instance in Ishizaka and Nemery (2013); here we focus in particular in decision making for lifecycle investment.

An example of a support system for decision making by pension fund boards is given by Dert and Leegwater (2011), who describe a procedure that is followed at ABN AMRO Pension Fund (AAPF). The selection process starts with a broad range of about forty strategies. An ALM study is performed to determine attributes of these strategies such as the expected funding ratio after one year and after fifteen years, the required level of contributions, the probability of underfunding over the next fifteen years, the potential for indexation of benefits, and so on. On the basis of this study, the investment advisory committee selects about ten strategies to be presented to the board of trustees. The board is informed about the attributes of the strategies as determined in the ALM study, but not about the way in which these strategies would be implemented (for instance, which asset mix would be used).

After the choice has been made, the implementation is revealed to the board, which then in principle can still change its mind (but typically doesn't).

The number of attributes that is presented to the board of trustees of AAPF is quite large; the information sheet shown as an example by Dert and Leegwater (2011) lists a total of nine attributes. For decision making by an individual, this number seems rather high. On the basis of the behavioral finance literature, we discuss in this paper how the dimensionality of the attribute space can be reduced to the extent that the final tradeoffs can be represented graphically in two-dimensional space. Inevitably, reduction of the attribute space makes preselection of strategies more important. We show the effects of preselection below in a simple test case by employing different classes of candidate strategies.

This paper contributes to the literature in two directions. The academic literature on the design of lifecycle strategies, in what might be called the Merton tradition, has placed emphasis on the optimization of a single objective, which, for tractability as well as for consistency, is often of the expected-utility form. We propose an alternative evaluation method, which articulates the quality of parametrized *classes* of strategies by means of the *tradeoffs* they generate between key attributes.¹ The traditional optimization methods are still valuable in this context, because they can be used as ways of producing classes of candidate strategies. Secondly, we extend the behavioral literature in the pensions field, in what might be called the Thaler tradition, by devising decision support for committed and well informed individuals. The behavioral literature is geared towards finding ways of supplying implicit support (“nudges”, in the terms of Thaler and Sunstein (2008)) to individuals of low or moderate financial literacy. However, anecdotal evidence suggests that even financially very literate persons have difficulties in arriving at lifecycle investment decisions, and in particular do not find it easy to specify their preferences in terms of utility functions. We turn to the behavioral literature to extract elements that can be used effectively within a decision process that is intended to be fully rational.

The organization of the remainder of the paper is as follows. In the following Section 2 we take cues from the behavioral literature in order to select key attributes around which a decision process on lifecycle investment can be organized. In the rest of the paper we develop the structure of a decision process and illustrate it in a small case study. Section 3 describes the setting of the case study. To show the effect of preselection, we use two classes of strategies which are presented in Section 4. Results of the case study follow in Section 5. Finally, Section 6 contains the conclusions.

2 Reduction of attribute dimensionality

The result of a lifecycle investment strategy may be described in terms of the probability distribution of the accumulated capital at retirement age. This is in fact a simplified description; additional elements may include for instance the psychological effects (either hopeful or fearful anticipation) that can be generated by strategies during the phase towards retirement, and nonprobabilistic forms of uncertainty (ambiguity) associated with projections over long time periods. Even in the relatively

¹The term “attribute” is used in this paper both for directly controlled (chosen) properties of a strategy, such as the contribution level, and for indirectly controlled (emerging) properties, such as the expected level of retirement income.

simple description form of a probability distribution, the outcome of the strategy decision is already an object that cannot be captured completely in terms of finitely many parameters. The question then arises how the dimension of the attribute space can be reduced in order to make an effective decision process possible.

A classical answer is to select a utility function. Indeed, once a utility function is chosen, then (in a complete market) the optimal probability distribution is obtained by applying the inverse marginal utility function to the pricing kernel times a multiplication factor that is determined by the budget constraint (Karatzas and Shreve, 1998). Viewed from a choice perspective, the method of expected utility optimizing replaces the problem of selecting an object with an infinite-dimensional attribute space (the probability distribution of available capital at retirement age) by the problem of selecting another object, which however is infinite-dimensional as well (the utility function). It may be argued that there is still progress here, since a utility function, once chosen, can be used not only for decisions on lifecycle investment, but also for other decisions. However, in view of the rather limited use of this methodology in practice, however, the argument does not appear strong. The question remains whether it is easier to choose a utility function to express one's preferences, or to choose a probability distribution directly.

Alternative preference specifications may be taken from the extensive literature on behavioral finance which has developed since the seminal paper of Kahneman and Tversky (1979). The orientation of this literature is primarily descriptive, so that care should be applied in using it for purposes of informed choice. In many experimental studies, subjects are confronted with stakes that are relatively small in comparison to lifecycle investment; they are likely to interpret the experimental environment as representative of choice situations that may reoccur more or less frequently, in contrast to saving for retirement. Nevertheless, the descriptive research indicates crucial elements in human decision making, leading to important pointers for normative theory.

A central element in the work of Kahneman and Tversky (1979) is the notion of a *benchmark*, by which outcomes can be classified as either “gains” or “losses”. A similar notion is present in the SP/A (security-potential/aspiration) theory of Lopes (1987) in the form of the *aspiration level*. Lopes describes the aspiration level as a situational variable (relating to the environment), but she also states that its choice may be influenced by the dispositional factor (relating to the individual) of security versus potential, which describes the extent to which a person is inclined to emphasize more the downside (security) or the upside (potential). The aspiration level is operationalized by Lopes and Oden (1999) as a probabilistic constraint, meaning that subjects choose their strategies in such a way that the probability of achieving the aspiration level is at least a given percentage (the *confidence level*). An analogous notion is formulated in the “utility mass” model of Manski (1988). The importance of the probabilities of ending up above or below a certain benchmark in human decision making has been confirmed in experimental studies (Fennema and Wakker, 1997; Payne, 2005). The attitude of businessmen in pondering possible projects was described by Shackle (1949) as one that takes into account, on the one hand, the best outcome that can reasonably be hoped for, and on the other hand, the worst outcome that is sufficiently likely so that it has to be taken into account. This description, although dating from the first half of the twentieth century, would still seem to be useful today.

The term “aspiration level” suggests a goal-type reference point, incorporating a certain amount of

ambition. The use of a benchmark at a low level, interpreted as representing a minimum standard, is also quite familiar in finance, the “safety first” model of Roy (1952) being an early example. In a study of managerial incentives, March and Shapira (1987) find that managers appear to work with two focal values rather than a single one; these values typically correspond to a target level for performance and a survival level. An analogous theory has been proposed in behavioral ecology by Hurly (2003) to explain the foraging behavior of hummingbirds under risk. The two benchmark levels in this case correspond to survival and to reproductive ability. A tri-reference-point theory of decision making under risk was proposed by Wang and Johnson (2012). The reference points correspond to minimum requirements, status quo, and goal. Empirical evidence for multiple reference points in human decision making is shown by Koop and Johnson (2012).

The SP/A theory was criticized by He and Zhou (2016) for combining elements representing “security” and “potential” in a simple way, by taking a convex combination. The approach advocated in He and Zhou (2016) instead rests on a more involved aggregation of these elements, expressed by means of an inverse-S-shaped distortion of probabilities. This approach is similar to the classical approach via expected utility in that it arrives at a single objective, which can then be used in an optimization procedure. Individuals may in principle express their dispositional characteristics by selecting a suitable distortion function, or presumably (since, again, the distortion function is an infinite-dimensional object) by choosing particular values of parameters within one of the available parametrized classes of distortion functions. One may in fact attempt to elicit such parameter values by experiments. This methodology is fraught with hazard, however, given the disputable reliability of the decision model across the wide range stretching from laboratory experiments to real-life decisions on investing for retirement. An advantage of multiple objectives is that one can avoid specifying an explicit model for the way in which these objectives are aggregated by the decision maker. Stated differently, this means that the decision maker is allowed to control multiple objectives directly, rather than indirectly by the choice of certain parameter values in a decision model.

In this paper, we take the presence of multiple reference points as a guideline for the design of a decision process. Within the context of lifecycle planning, it would seem reasonable not to distinguish between “status quo” and “goal”, since the purpose of retirement planning is often stated as maintaining the standard of living. We shall therefore work below with two focal points, interpreted respectively as an aspiration level and a subsistence level. In the style of Lopes and Oden (1999), the probability of achieving the aspiration level will be specified as an input.

3 Setting for case study

Since the focus of the present paper is on the design of a decision framework rather than on economic modeling, and because we want to limit the size of computations that express relevant attributes in terms of chosen strategy classes, we will work with a simple model. Obviously, in actual applications, the model should be replaced by one that is more realistic.

3.1 Situational variables: economic environment

The purpose of this section is to construct an environment in which the idea of a multi-objective decision framework for lifecycle investment can be illustrated. For reasons mentioned above, we use a simple model, namely the standard Black-Scholes model. The only risk factor is stock market risk, traded through a stock index S_t driven by Brownian motion. The riskless bond B_t is subject to exponential growth at constant rate r . The asset price processes are defined by

$$\begin{aligned} dS_t &= \mu S_t dt + \sigma S_t dW_t \\ dB_t &= r B_t dt \end{aligned}$$

where μ and σ are the constant drift and volatility parameters, and W_t is a standard Brownian motion. We assume that the prices of all assets are expressed in terms of an inflation-indexed numéraire, so that effectively the rate of inflation as described in the economy above is zero. In the numerical examples in later sections, the following parameter values are assumed: real riskless interest rate $r = 2\%$, expected real stock market return $\mu = 5\%$, and stock market volatility $\sigma = 18\%$. This model ignores many issues including the presence of multiple asset classes, variability of interest rates, and market frictions; it is meant for illustrative purposes only.

We consider an individual who starts her active career at age 25 and retires at 65. During the active period the individual works and earns a flat real labor income Y . In other words, labor income is taken to be riskless, and wage inflation is equal to price inflation. A fixed part of the labor income is saved for consumption after retirement, and this is the only source of retirement income. At age 65, the available capital is converted to a fixed (that is, inflation-indexed) annuity, at a conversion rate that is given by the simplified formula

$$\nu = \int_0^{T_D - T} e^{-rs} ds$$

where T_D denotes the expected time of death. In calculations below, we take $T_D - T = 20$. Again, these assumptions imply strong simplifications.

In the context of the economy above, a *lifecycle strategy* consists of a choice of the contribution rate and a choice of an investment strategy for retirement savings during the accumulation phase. While the setting that is constructed in this way is rather simple, it does represent some of the major choices that need to be made in retirement planning.

3.2 Dispositional variables: key attributes

The key attributes that we propose relate to two levels of retirement income, which we refer to as *desired income* (target income) and *minimum required income* (subsistence income). These are attributes that are obtained as results of any given strategy. The minimum required income that we associate with a given strategy is the income that is achieved by this strategy with a given (high) probability $1 - p_1$. Likewise, the desired income associated with a strategy is defined as the income that is achieved by the strategy with a second (lower) probability p_2 . More precisely, the definition that we use is the following.

A given lifecycle strategy, consisting of a contribution rate and an investment strategy, produces a (usually random) portfolio value at the time of retirement T . This portfolio value will be denoted

by A_T . The minimum required income κ_1 that is associated with the strategy is defined, analogously to Value-at-Risk, as the p_1 -quantile of the distribution of retirement income:

$$\kappa_1 = \sup \{ x \mid \mathbb{P}(\nu^{-1}A_T < x) \leq p_1 \}. \quad (1)$$

Likewise, the desired income κ_2 associated with the given strategy is defined as the p_2 -quantile of the distribution of retirement income that is generated by the strategy:

$$\kappa_2 = \sup \{ x \mid \mathbb{P}(\nu^{-1}A_T \geq x) \geq p_2 \}. \quad (2)$$

Under the assumption that the distribution of portfolio value A_T is continuous, the two income levels can also be obtained as the solutions of the implicit equations

$$p_1 = \mathbb{P}(\nu^{-1}A_T \leq \kappa_1), \quad p_2 = \mathbb{P}(\nu^{-1}A_T \geq \kappa_2). \quad (3)$$

The *contribution rate* is obviously an important attribute as well. Tradeoff diagrams to be constructed below will include the impact of the contribution rate.

4 Strategies

In this section, we define two classes of strategies, which will subsequently be compared in terms of the tradeoffs that they generate. For the same reasons as mentioned at the beginning of the previous section, the strategies that will be considered here are deliberately kept simple. More sophisticated strategies would be called for in actual applications.

4.1 Glide path strategies

As a first class of strategies to be considered, we take a collection of *glide path strategies*. Such strategies are often used in savings products offered by insurers. These strategies rebalance the equity allocation as a deterministic function of age only. Specifically, we consider strategies that initially allocate 100% of assets to stocks and that shift to the riskless portfolio at retirement. The speed at which different strategies reduce the equity exposure varies over the lifecycle. Figure 1 displays the age-related stock exposure α_t corresponding to the strategies that we consider.

Insurance companies that offer glide path strategies often do not present a wide range of such strategies to their customers, but typically limit the choice to three options (for instance called “aggressive”, “defensive”, and “neutral”). We highlight three strategies in Fig. 1 and refer to them as *risk averse*, *linear*, and *risk seeking*. These schemes respectively reduce the equity exposure early in life, linearly, or only later in life. The performance of these schemes will be discussed more extensively below.

4.2 Collar strategies

4.2.1 Construction

In addition to the glide path strategies discussed above, we consider in this paper a class of *collar strategies*. These are strategies that replicate a payoff as shown in Fig. 2. There are four parameters,

Figure 1: Glide path strategies.

namely a low strike price K_1 , a high strike price K_2 , a guaranteed payoff level (floor) θ_1 , and a maximal payoff level (cap) θ_2 . The payoff can be constructed by a combination of a long and a short position in call options with different strikes, as shown in Fig. 2 as well. The payoff can also be considered as resulting from a long position in risky assets together with a protective put and a written call which helps reduce the amount needed to set up the replicating portfolio. An expression for the payoff is

$$V_T = \max \left\{ \theta_1, \min \left\{ \theta_1 + \frac{\theta_2 - \theta_1}{K_2 - K_1} (S_T - K_1), \theta_2 \right\} \right\}. \quad (4)$$

Alternatively, the payoff can be written as

$$V_T = \theta_1 + a (F(S_T, K_1) - F(S_T, K_2)) \quad (5)$$

where $a = \frac{\theta_2 - \theta_1}{K_2 - K_1}$ represents the slope of the linearly increasing part of the option payoff, and $F(S_T, K) = \max\{S_T - K, 0\}$ denotes the payoff function of a call option.

Figure 2: Option composition of the collar strategy.

The strike prices K_1 and K_2 relate to the likelihood of realizing the cap and the floor, respectively. In order to reduce the dimension of the parameter space for the class of collar strategies, we fix the strike prices K_1 and K_2 for given levels of the floor and the cap by prescribing the probabilities by which these levels should be achieved. These probabilities will be denoted by p_1 and p_2 , respectively; numerical values that will be used below are $p_1 = 2.5\%$ and $p_2 = 70\%$. The lower and upper strike are then solved from

$$p_1 = \mathbb{P}(S_T \leq K_1) = 1 - \Phi\left(\frac{\log(S_0/K_1) + (\mu - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right)$$

$$p_2 = \mathbb{P}(S_T \geq K_2) = \Phi\left(\frac{\log(S_0/K_2) + (\mu - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right)$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. Note that the parameter μ is used here, rather than r , since we assume that the investor is interested in the real-world probabilities of reaching the cap and the floor, rather than the probabilities under the risk-neutral measure.

4.2.2 Implementation

In a complete market, such as the Black-Scholes market, a contingent claim can be replicated by suitable trading in the underlying asset. The initial wealth required to finance the replicating portfolio is equal to the unique price of the claim. The replication strategy itself is uniquely defined as well; at each point in time, the amount of risky assets to be held in the portfolio is determined by the delta, which is the first derivative of the current option value with respect to the current value of the underlying asset. Any portfolio value remaining after the required investment in stocks is invested in the riskless bond. If the replicating portfolio is rebalanced such that the number of stocks in the portfolio is equal to the delta at any instant, then the payoff is achieved perfectly at retirement. Since the collar payoff can be written as a linear combination of call option payoffs, both the price and the replicating strategy can be worked out analytically.

The investment strategy is defined by the aggregate of strategies that replicate the floor, the long call and the short call. Whereas the minimum income constraint is maintained by bond investments only, the call options are replicated by dynamic portfolios of both bonds and stocks. If the present value of the payoff approximates the discounted value of the floor, stock exposure tends to zero because all assets have to be invested in the bond to maintain the minimum income constraint. If, on the other hand, the present value of the payoff approximates the discounted value of the cap, stock market risk is reduced as it is no longer needed to achieve the desired income. The sensitivity of the investment strategy to the stock price is never exactly zero in the model that we use, but it is clearly the largest in situations when the option value lies between the discounted values of the cap and the floor (cf. Fig. 3).

To illustrate the asset allocation, we assume that the minimum required replacement rate κ_1 and the desired replacement rate κ_2 are equal to 50% and 80% of wage, respectively.² The retirement age

²The concept of replacement rate is widely used by pension professionals in setting an aspiration for pension plans. It is also common in academic research as a measure for well-being after retirement (Haveman et al., 2007; Knoef et al., 2016; OECD, 2015). The chosen percentages can be appropriate if the agent will have additional sources of income during retirement, such as a state pension.

is fixed at 65, and the quantiles corresponding to desired income and minimum income are 70% and 2.5%, respectively. The contribution rate (as computed by the formula in (9) below) equals 17.5%.

Figure 3 illustrates the fraction of accumulated savings invested in stocks as a function of time and annualized real return. In the first years of saving and investing, the strategy is short in the riskless bond.³ The position is short in the bond because the discounted value of future contributions exceeds the amount of bond investments required to replicate the payoff. In other words, the stock of human capital is large enough to construct the bond positions required for replicating the floor and the call options by using anticipated future contributions only. As a result, there is a leveraged position in stocks to improve the likelihood of achieving the desired income. Savings will be allocated to bonds only once human capital (i.e., the discounted value of future contributions) becomes insufficient.

Figure 3: Fraction of accumulated savings invested in stock index.

Also in Bodie et al. (1992), young individuals borrow at the risk-free rate to finance investments in the risky asset. However, by maximizing expected power utility, an individual typically wants to borrow a larger amount to acquire the optimal asset allocation. In the presence of riskless labor income, Bodie et al. (1992) show that a young individual may want to borrow as much as five times her annual wage to invest in stocks.⁴ In order to construct the collar payoff, an individual at the start of her career would have to borrow 75% of her annual wage. In more intricate models, in which possible dependence between labor income and stock index returns is taken into account (for instance Benzoni et al. (2007)), the leverage may be lower.

In practice, borrowing against the riskless asset may be hard to realize because young participants face borrowing constraints. If the individual takes part in a collective fund, then the desired exposure may be realized more easily. The fund can implement an aggregate investment policy based on desired exposures of all participants, young and old. The averaging would lead to a limited level of exposure, and moreover would dampen the volatility in the investment mix that may arise in a collar strategy

³Leveraged equity exposure is not displayed in the figure, in order not to compromise visibility of other parts of the graphic.

⁴Given power utility, the optimal fraction of total capital (accumulated savings plus discounted future contributions) invested in stocks is given by $\frac{\mu-r}{\gamma\sigma^2}$, where γ denotes the risk aversion parameter. Values for γ in the range of 3 to 5 are often considered typical.

when carried out by an individual (see Fig. 3). By age-dependent allocation of the investment returns, the fund can realize different risk profiles for different groups of participants.

4.2.3 Calculation of key attributes

The contribution rate is determined by the market price of the portfolio required to replicate the payoff. As a complete Black-Scholes market is assumed, risk-neutral pricing techniques can be used to determine the unique price of the contingent claim.

Let $\mathbb{Q}$ be the unique equivalent martingale measure corresponding to the bond as numéraire. According to the fundamental theorem of asset pricing, the process defined by $B_t^{-1}V_t$ is a martingale under the equivalent measure $\mathbb{Q}$ and therefore the value of the option at time $t \leq T$ satisfies

$$V_t = E_t^{\mathbb{Q}} \left[\frac{B_T}{B_t} V_T \right] = e^{-r(T-t)} E_t^{\mathbb{Q}} V_T. \quad (6)$$

Solving for the option value at time t yields a function depending on time t and the value on the underlying stock market index S_t at time t . The option value at time t is given by

$$V_t = aS_t(\Phi(d_1) - \Phi(d_2)) + e^{-r(T-t)} \left(\theta_1 - a(K_1\Phi(d_1 - \sigma\sqrt{T-t}) - K_2\Phi(d_2 - \sigma\sqrt{T-t})) \right) \quad (7)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function,

$$d_i = \frac{\log(\frac{S_t}{K_i}) + (r + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}, \quad i \in \{1, 2\} \quad (8)$$

and $a = \frac{\theta_2 - \theta_1}{K_2 - K_1}$. Upon entry into the pension plan, the value of the payoff is equal to V_{t_0} .

The contribution rate that is associated with a given payoff profile can be computed from the principle that, at the moment of entry into the pension plan, the sum of the discounted contribution payments must be equal to the price of the payoff. This is sufficient to guarantee that the replication strategy is self-financing. Given the assumption that participants contribute a constant fraction of labor income over the active period of the lifecycle, the contribution rate c is therefore solved from

$$V_{t_0} = cY \int_0^T e^{-rs} ds. \quad (9)$$

By our assumption that prices (including the wage Y) are given relative to an inflation-indexed numéraire, the discount rate equals the real riskless rate r .

5 Tradeoffs

By fixing the retirement age and the strike prices, the dimensionality of the saving and investment decision reduces to choosing three variables: contribution, desired income, and minimum required income. These parameters represent the embedded tradeoff between consumption during working life and consumption in retirement. Given the contribution rate that an individual can afford to save, she is in a position to select a feasible combination of desired income and minimum required income.

Figure 4 shows a number of tradeoff curves, labeled by the corresponding contribution rates. These curves quantify the relation between contribution, desired income, and minimum required income. For the fixed parameters considered, the following statements can be made.

Figure 4: Trade-off in terms of desired and minimum required income.

- Given the contribution rate, a 1 percentage point increase of the minimum required income requires a 1 percentage point decrease of the desired income.
- For 1 percentage point additional contribution, the desired income, the minimum required income, or a combination of both may be increased by 7.4 percentage points on aggregate.

Note that the slope of the tradeoff curve depends on the strike prices. *Ceteris paribus*, the slope is decreasing in K_2 and increasing in K_1 . For example, if the probability of achieving the desired income increases, i.e. the upper strike K_2 decreases, the slope of the tradeoff curve flattens.

In addition to the three attributes that are represented Figure 4, a decision maker may want to consider other attributes as well. For instance, an individual may want to think about the possibility of retiring earlier or later than the age of 65 that has been supposed above, or she may wish a different probability of realizing the target. Under such alternatives, similar tradeoff curves as shown in Figure 4 can be constructed. Retiring early makes a given combination of desired and minimum required income more expensive because the investment period shortens, total contributions decrease, and the remaining life expectancy at retirement increases. Hence, the tradeoff curves shift in the direction of the origin. Increasing the probability of achieving the desired income also requires more contribution for a given combination of desired and minimum required income. When the contribution rate is kept unchanged, the tradeoff curve rotates around the point that corresponds to the riskless portfolio. The tradeoff curve flattens, and therefore increasing the desired income becomes more expensive in terms of the minimum required income.

Figure 5 compares the tradeoff curve for the collar strategy proposed in this paper to the tradeoff curve for the lifecycle strategies. The contribution rate has been set to 17.5%. It is seen that, given a minimum required level of income replacement, the collar strategies provide a larger desired income than the lifecycle strategies do. The collar strategy is capable of yielding levels of desired income that are unattainable by implementing a lifecycle strategy. In order to achieve similar desired income levels for the class of lifecycle strategies, an agent would have to pay considerably more contribution. Relative to the collar strategy, households would therefore save too much and face the regret of forgone

consumption opportunities during the many years before retirement.

Figure 5: Tradeoff between desired and minimum required income.

6 Conclusion

In a large part of the literature on behavioral economics, attention is focused on individuals who make “mistakes” due to “behavioral biases”, and who therefore must be “nudged” towards decisions which the nudge designer knows are better. While the need for studies of this type is undeniable, there are also many situations in which individuals who are intelligent, responsible, and well-informed still have difficulties in reaching a good decision. The aim of this paper has been to contribute to what might be called the architecture of informed choice, with focus in particular on the lifecycle investment decision.

In the academic literature on lifecycle investment, the predominant approach has been to construct first a decision model that arrives at a single criterion by which alternative strategies can be ranked, and then to optimize this criterion under certain assumptions concerning the economic environment. Characteristics that differ between individuals, such as attitude towards risk, can be expressed by choosing different parameters in the decision model. The degree of acceptance of this approach in practice has been limited; decision makers do not find it easy to express their preferences in terms of the parameters of the typical decision models that are used. In this paper, an alternative approach has been suggested, based on less-than-complete aggregation of preferences. The approach calls for identification of a limited set of key attributes and for the construction of tradeoff diagrams between these attributes, such as shown in Fig. 4 (right panel).

The proposed decision process is based on the selection of a set of key attributes, as well as on the selection of a class of strategies. Both choices involve a certain amount of nudging, used here to provide support to knowledgeable decisionmakers rather than to correct behavioral biases. Given the size of the attribute space that in principle can be taken into account and the limited information processing capabilities of even the most well trained humans, it seems inevitable that the freedom of choice should be limited. The main issue is whether the imposed limitations do not eliminate any opportunities that are considerably more attractive than the ones that are still left open.

As shown in the example discussed in this paper, it may happen that one class of strategies uniformly dominates another class with respect to a chosen set of key attributes. Uniform dominance induces a partial preference relation between classes of strategies. Candidate classes of strategies can be chosen on an ad hoc basis, as in the example in this paper. Alternatively, they may be obtained by optimization with respect to a parametrized decision model, such as expected utility with a parametrized class of utility functions; different choices of parameters lead to different strategies. One might also think of devising optimal classes of strategies that are in a sense optimal, by optimizing with respect to weighted linear combinations of the key attributes. One reason not to follow the latter approach may be that the resulting optimization problems are difficult to solve. Another reason could be that such an approach ignores attributes that are not listed among the key attributes, but that may still not be completely negligible, or attributes that are in fact important, but that are difficult to quantify. Focusing exclusively on the designated key attributes may lead to solutions that may be optimal (in a restricted sense) but are actually undesirable. The constraint imposed by a restricted class of strategies may serve to protect other attributes than the ones that appear in the tradeoff diagram, and may therefore constitute a form of nudging as already alluded to above. The task of the financial engineer in this context could be stated as the design of classes of strategies that provide good tradeoffs between key attributes and that are sensible from the point of view of protection of secondary attributes and attributes that are hard to quantify.

Obviously, the economic model used in this paper is meant for illustrative purposes only. For the purpose of actual decision making on lifecycle investment strategies, more elaborate ALM models are needed, which take into account for instance the variability of interest rates and the uncertain evolution of mortality tables. A more complete list of attributes to be considered should be used. Further work on the architecture of informed choice might be supported by laboratory experiments and the practical experience of advisors.

The analysis in this paper has mainly focused on the case of individual lifecycle planning, although it has been noted that collective funds can be used to alleviate credit constraints. The board of trustees of a collective fund could follow a similar path to decisions as has been suggested above for individuals, by identifying key attributes of possible strategies and considering the tradeoffs between them. In a more elaborate approach, the board might place itself in the position of various groups of participants (for instance, different ages) and work out an investment policy for these groups, finally arriving at its own investment policy by summation of the investment decisions that the board assumes would be taken by the participants themselves. This also means that the investment returns would be allocated accordingly. A third possibility is that the board decides on a choice palette that will be offered to plan members, for instance in the form of a tradeoff diagram, from which participants can choose; subsequently, the board may decide on its investment policy based on the choices of the plan members. The options that are open here call for further investigation.

Another possible avenue of future research relates to the use of multiple objectives in dynamic optimization. The classical approach to lifecycle investment, as in the work of Merton (1969, 1971) and in many other papers, makes use of dynamic programming. The central object in this approach is the value function, which is a single-valued function that is designed to aggregate all of the decision maker's objectives. Aggregation takes place with respect to objectives as well as with respect to uncertainty. This is not compatible with lines of thinking in which different objectives relate to uncertainty in

different ways. Mean-variance analysis is a classical example of such a line of thinking, and indeed it is well known that it is not straightforward to merge MV analysis with dynamic programming. However, a combination has been proposed by Basak and Chabakauri (2010), on the basis of a backward recursion as in standard dynamic programming, but using two quantities in the recursion, rather than one. In this way, a multiple-objective element is retained in a dynamic programming environment, ensuring time consistency. It would be of interest to see whether this idea can be extended for instance to objectives expressed in terms of quantiles, as in this paper.

References

- Suleyman Basak and Georgy Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23:2970–3016, 2010.
- Shlomo Benartzi and Richard H. Thaler. Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives*, 21:81–104, 2007.
- Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein. Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *Journal of Finance*, 62:2123–2167, 2007.
- Zvi Bodie and Henriëtte M. Prast. Rational pensions for irrational people: behavioral science lessons for the Netherlands. In Lans Bovenberg, Casper van Ewijk, and Ed Westerhout, editors, *The Future of Multi-Pillar Pensions*, pages 299–329. Cambridge University Press, Cambridge, 2012.
- Zvi Bodie, Robert C. Merton, and William F. Samuelson. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control*, 16:427–449, 1992.
- Gordon L. Clark, Emiko Caerlewy-Smith, and John C. Marshall. Pension fund trustee competence: decision making in problems relevant to investment practice. *Journal of Pension Economics and Finance*, 5:91–110, 2006.
- Cees L. Dert and Geraldine Leegwater. Decision-making and solvency-based dynamic asset allocation at the ABN AMRO Pension Fund. *Journal of Investment Consulting*, 12(2):35–41, 2011.
- Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75:643–669, 1961.
- Barbara Fasolo, Gary H. McClelland, and Peter M. Todd. Escaping the tyranny of choice: When fewer attributes make choice easier. *Marketing Theory*, 7:13–26, 2007.
- Hein Fennema and Peter Wakker. Original and cumulative prospect theory: A discussion of empirical differences. *Journal of Behavioral Decision Making*, 10:53–64, 1997.
- Robert Haveman, Karen Holden, Andrei Romanov, and Barbara Wolfe. Assessing the maintenance of savings sufficiency over the first decade of retirement. *International Tax and Public Finance*, 14:481–502, 2007.
- Xue Dong He and Xun Yu Zhou. Hope, fear, and aspirations. *Mathematical Finance*, 26:3–50, 2016.

- Andrew T. Hurly. The twin threshold model: risk-intermediate foraging by rufous hummingbirds, *Selasphorus rufus*. *Animal Behaviour*, 66:751–761, 2003.
- Alessio Ishizaka and Philippe Nemery. *Multi-criteria decision analysis: methods and software*. Wiley, Chichester, 2013.
- Sheena S. Iyengar, Gur Huberman, and Wei Jiang. How much choice is too much? Contributions to 401(k) retirement plans. In *Pension Design and Structure: New Lessons from Behavioral Finance*, pages 83–95. Oxford University Press, Oxford, 2004.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–291, 1979.
- Ioannis Karatzas and Steven E. Shreve. *Methods of Mathematical Finance*. Springer, New York, 1998.
- Marika Knoef, Jim Been, Rob Alessie, Koen Caminada, Kees Goudswaard, and Adriaan Kalwij. Measuring retirement savings adequacy: developing a multi-pillar approach in the Netherlands. *Journal of Pension Economics and Finance*, 15:55–89, 2016.
- Gregory J. Koop and Joseph G. Johnson. The use of multiple reference points in risky decision making. *Journal of Behavioral Decision Making*, 25:49–62, 2012.
- Lola L. Lopes. Between hope and fear: The psychology of risk. *Advances in Experimental Social Psychology*, 20:255–295, 1987.
- Lola L. Lopes and Gregg C. Oden. The role of aspiration level in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology*, 43:286–313, 1999.
- Annamaria Lusardi and Olivia S. Mitchell. Financial literacy and retirement planning in the United States. *Journal of Pension Economics and Finance*, 10:509–525, 2011.
- Charles F. Manski. Ordinal utility models of decision making under uncertainty. *Theory and Decision*, 25:79–104, 1988.
- James G. March and Zur Shapira. Managerial perspectives on risk and risk taking. *Management Science*, 33:1404–1418, 1987.
- Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, pages 247–257, 1969.
- Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3:373–413, 1971.
- Olivia S. Mitchell and Stephen P. Utkus, editors. *Pension Design and Structure: New Lessons from Behavioral Finance*. Oxford University Press, 2004.
- OECD. *Pensions at a Glance 2015. OECD and G20 Indicators*. OECD Publishing, Paris, 2015.

- John W. Payne. It is whether you win or lose: The importance of the overall probabilities of winning or losing in risky choice. *Journal of Risk and Uncertainty*, 30:5–19, 2005.
- Andrew D. Roy. Safety first and the holding of assets. *Econometrica*, pages 431–449, 1952.
- Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- George L.S. Shackle. *Expectation in Economics*. Cambridge University Press, Cambridge, 1949.
- Richard H. Thaler and Shlomo Benartzi. Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112:164–187, 2004.
- Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press, New Haven, 2008.
- X.T. Wang and Joseph G. Johnson. A tri-reference point theory of decision making under risk. *Journal of Experimental Psychology: General*, 141:743–756, 2012.