

Dynamic Optimization for Multi-Goals Wealth Management *

Sanjiv R. Das
Santa Clara University

Daniel Ostrov
Santa Clara University

Anand Radhakrishnan
Franklin Templeton Investments

Deep Srivastav
Franklin Templeton Investments

May 25, 2021

Abstract

We develop a dynamic programming methodology that seeks to maximize investor outcomes over multiple, potentially competing goals (such as upgrading a home, paying college tuition, or maintaining an income stream in retirement), even when financial resources are limited. Unlike Monte Carlo approaches currently in wide use in the wealth management industry, our approach uses investor preferences to dynamically make the optimal determination for fulfilling or not fulfilling each goal and for selecting the investor's investment portfolio. This can be computed quickly, even for numerous investor goals spread over different or concurrent time periods, where each goal may be all-or-nothing or may allow for partial fulfillment. The probabilities of attaining each (full or partial) goal under the optimal scenario are also computed, so the investor can ensure the algorithm accurately reflects their preference for the relative importance of each of their goals. This approach vastly outperforms buy and hold strategies and target-date funds, widely used in the wealth management industry.

Keywords: Goals-based investing, dynamic programming, retirement planning.

JEL Codes: G11, G40, G50

*We thank the editor and two referees for many constructive comments. We are grateful for discussions and contributions from many of the team at Franklin Templeton Investments, especially Aviva Casanova for extensive input, and Jennifer Ball for her thoughts and support on this research. Thanks also to Hersh Shefrin for his thoughtful feedback. We received feedback from many seminar participants at Amazon Web Services, Indian Institute of Management (Kozikhode), Florida International University, QWAFEFW CFA San Francisco, Pacific Institute for Mathematical Sciences at the University of Manitoba, Fields Institute Toronto, Behavioral Advisor Conference, San Francisco, Journal of Investment Management Conference Boston, The Free University Amsterdam, De Nederlands Bank, Amsterdam. Authors: Sanjiv Das (srdas@scu.edu, 408-554-2776) and Dan Ostrov (dostrov@scu.edu, 408-554-4551) are at 500 El Camino Real, Santa Clara, CA 95053. Anand Radhakrishnan (andy.radhakrishnan@franklintempleton.com, 800-632-2301) and Deep Srivastav (deepratna.srivastav@franklintempleton.com, 800-632-2301) are at 1 Franklin Pkwy #970, San Mateo, CA 94403.

1 Introduction

Goals-based wealth management (GBWM) is an investment philosophy focused on attaining the desired goal or goals specified by an investor (Chhabra (2005); Nevins (2004); Browne (1997); Browne (1999); Brunel (2015); Pakizer (2017); Parker (2020); Parker (2021)). This paper’s contributions to GBWM include presenting a dynamic programming algorithm that enables an investor with multiple purchasing goals to quickly determine how to optimally invest their money and optimally prioritize which goals they should or should not purchase — even if there are hundreds of goals over a long time horizon. Further, we show that this algorithm’s solution substantially outperforms current Monte Carlo financial planning approaches.

This investor based viewpoint of GBWM corresponds to a new notion of risk. Traditionally, risk is defined as the volatility of the investments in an investor’s portfolio. In contrast, for GBWM with a single goal, risk is defined as the probability that an investor does not meet that financial goal. So, for example, if a young person moves all their money from stock into cash for the goal of retirement, they are decreasing their risk from a traditional point of view, but increasing their risk from the GBWM point of view.

For multiple goals in well-funded portfolios, the GBWM notion of risk for a single goal translates easily to multiple goals. It becomes the probability that the investor will not be able to attain all their goals (relevant to lifestyle risk as in Bergerson et al. (2016)). But of considerably more importance is the complex question of what, optimally, should be done when limited investor resources necessitate prioritizing and then choosing among the investor’s multiple competing goals (Consiglio et al., 2004), as we investigate in this paper, especially over long horizons, through retirement (Simsek et al., 2018), or Kim et al. (2020) for a solution using stochastic programming. Our approach enables a flexible goal prioritization, while not mandating a strict sequence in which goals should be realized.

Consider, for example, the simple case where an investor has just two goals: Let’s say that in 5 years, the investor wants to take a nice vacation and in 10 years the investor wants to pay for their child’s first car. If they don’t have the money in their portfolio for the vacation at the end of five years, they must forgo the vacation goal of course. If they have a lot of money, they fulfill the vacation goal knowing that they will almost certainly be able to afford the car goal as well. But what if the investor has a moderate amount of money in their portfolio? Do they take the vacation? It depends not only on the cost of the vacation and the car, but also on the relative importance of the two goals to the investor. If they have just enough money to afford the vacation, but attaining the car goal, while more important, is likely out of reach due to being significantly more expensive, they should fulfill the vacation goal. If they have more money, they should switch to forgoing the vacation so as to optimize the chance of attaining the more important car goal, but if they have even more money they should again switch to taking the vacation because they can likely attain both goals. This leads to an obvious, key question: what values of portfolio wealth in 5 years optimally correspond to these switches regarding the decision to fulfill or not fulfill the vacation goal?

This simple two goal example quickly gets far more complicated when made more realistic. What if we consider partial goals like less expensive vacation choices instead

of the full goal of the nice vacation, or we consider the possibility of less expensive cars? What if we add additional goals like paying off a mortgage in 15 years or remodeling a kitchen in 5 years, concurrent with considering the vacation? Or, after 17 years, annually removing money from the portfolio for the 30 years after that to fund retirement? Optimally determining which goals to fulfill is not obvious, nor is it obvious what the optimal portfolio is at each time period to attain these goals. And if you add a projected income stream for the investment portfolio, how would that change these optimal decisions?

There is an extensive literature on dynamic programming approaches, using stochastic programming for asset management, such as the elegant works by [Mulvey and Vladimirou \(1992\)](#); [Dantzig and Infanger \(1993\)](#); [Consigli and Dempster \(1998\)](#); [Dempster, Germano, Medova, and Villaverde \(2003\)](#); [Topaloglou, Vladimirou, and Zenios \(2008\)](#). Specific to long-horizon portfolio planning for financial planning problems of individual investors, see [Mulvey and Shetty \(2004\)](#); [Consiglio, Cocco, and Zenios \(2004\)](#); [Consiglio, Cocco, and Zenios \(2007\)](#), and [Infanger \(2008\)](#). For managing individuals' portfolios in defined contribution plans, see [Konicz and Mulvey \(2015\)](#). See [Dempster and Medova \(2011\)](#) for individual planning problems as asset-liability problems, and [Konicz, Pisinger, Rasmussen, and Steffensen \(2015\)](#) for personal finance and pensions; [Konicz et al. \(2016\)](#) for the specific planning problem of a couple over time with uncertain mortality and including the complexity of inflation-linked annuities ([Konicz and Mulvey, 2013](#); [Konicz, Pisinger, and Weissensteiner, 2015](#)). We complement these large state-space problems to handle a large number of optional multiple goals in an efficient manner. The remainder of this introduction highlights the features of our approach that extend this literature in the framework of goals-based wealth management.

For multiple non-stochastic goals, we employ dynamic programming, which evolves backwards in time, to determine the optimal investment portfolio for an investor and, simultaneously, optimally determine for each goal if the investor is best off fulfilling the full goal, fulfilling a partial goal, or forgoing the goal completely. This optimization computation is usually completed in under 2 seconds. Even in the most complicated case we consider — which is in [Section 4.5](#), where we look to optimize a case that has a 60 year portfolio horizon, over 1000 potential wealth values, 15 investment portfolio choices at each time and wealth value, and 301 competing full goals with 138 partial alternative goals — the optimal solution is computed in only 17 seconds on a basic desktop computer by exploiting the ability to optimize the investment selection separately from the optimization of any goal fulfillment decisions. That is, the algorithm enables quickly computing which of the investor's goals should and should not be fulfilled optimally, as well as the optimal investment portfolio for the investor, at every time period given their portfolio's worth at that time.

The GBWM model works in harmony with both modern/rational portfolio theory and behavioral finance. Modern portfolio theory ([Markowitz \(1952\)](#); [Szegö \(2014\)](#)) prescribes the static portfolio that optimizes expected return for a fixed level of volatility. The set of these optimal portfolios forms the efficient frontier. Staying on the efficient frontier minimizes traditional risk, since it minimizes volatility for a given expected return. Staying on the efficient frontier also minimizes GBWM risk since, for a fixed level of volatility, we want the optimum expected return to attain goals. We therefore, ideally, only consider portfolios on the efficient frontier, although the algorithm will also work just as effectively if we are restricted to a set of portfolios that are not on the frontier.

This combination of mean-variance optimality and goals-based optimization addresses the criticisms of modern portfolio theory detailed in [Muralidhar \(2018\)](#).

Markowitz’s result for making optimal portfolio allocations was extended to dynamic models via maximizing the utility of the investor’s final wealth ([Merton \(1969\)](#); [Merton \(1971\)](#)). The utility of a wealth value corresponds to its importance or use to the investor. The primary roadblock in implementing this approach in practice is determining an appropriate utility function, which is investor specific. One problem in trying to determine an appropriate utility function is the fact that investors’ behavioral preferences can be in conflict with traditional utility function theory, as shown, for example, in Prospect Theory ([Kahneman and Tversky \(1979\)](#)). Behavioral finance work on the portfolio optimization problem has looked at how to appropriately embed behavioral considerations into the optimization. (See, for example, the papers by [Shefrin and Statman \(2000\)](#), [Das et al. \(2010\)](#), [Wang et al. \(2011\)](#), [Deguest et al. \(2015\)](#), [Alexander et al. \(2017\)](#), or the book by [Shefrin \(2008\)](#).) A second problem with determining an appropriate utility function is that there is no obvious feedback loop that enables the investor to gauge if a given utility function poorly fits their preferences and then enables the investor to correct the utility function. A third issue is that traditional utility functions equate utility with the cost of fulfilling the goal, whereas investors may have different priorities for goals that have the same cost, even at the same point in time.

The notion of utility is key in this paper because prioritizing goals requires that the investor specify the importance to them of each full or partial goal (i.e., each goal’s utility), in addition to determining each goal’s cost. Fortunately, in the context of this paper, assigning an appropriate utility value to each goal is relatively easy. In part, this is because we only need to assign utility values to a finite number of goals, as opposed to trying to specify a utility function over a continuum of wealth values. But also, we can take advantage of the nature and structure of the multiple goals scenario to address the two problems discussed above, namely, behavioral concerns and the need for a correcting feedback loop.

To do this, we initially create a potentially very coarse approximation for the utility values assigned to each goal and partial goal. The algorithm optimizes the expected value of the sum of the utilities from fulfilled goals and then uses the optimizing strategy to compute the corresponding probabilities of attaining each full or partial goal. The output of the algorithm is a table of goals with the optimized probabilities of achieving each goal. The fact the output is in terms of probabilities, as opposed to expected utility, is key to effectively communicating with investors.

Behavioral research shows that investors understand the notion of the probability of attaining a goal far better than most financial terms commonly used by wealth managers. See, for example, [Das et al. \(2018\)](#). This means that while investors are unlikely to understand the direct meaning of a utility assignment — for example, the difference between a utility of 200 versus 250 for a car goal is far from clear — they can understand the ramifications of their utility assignments to various goals by looking at the corresponding optimal goal probabilities. They can then alter their assigned utility values as finely as they wish to make these probabilities for attaining their goals conform to their preferences.

More specifically, if the investor finds the probability of attaining a specific goal too

low, they can just increase the utility assigned to that goal, with the understanding that it will lower the probability of attaining most of the other goals. Should they wish to increase the probability of attaining all of their goals, they can decide to add cash infusions into the portfolio at any specific time period or collection of time periods, or they can expand the range of investment portfolio strategies available to them. The investor can then use the algorithm to quickly recompute the new optimized probabilities for attaining each goal. This procedure may be iterated as many times as desired. This human-in-the-loop process enables us to directly determine the investor’s true preferences among their goals, instead of guessing or assuming them. In particular, it has the considerable advantage of not requiring the use of traditional utility functions over continuous domains that, by their nature, must be chosen in an ad-hoc manner that looks to approximate an investor’s preferences.

This approach yields the correct extension to understanding GBWM risk in the context of multiple competing goals. It is no longer a single probability. It is now a collection of probabilities that compete with each other, and this overall risk is minimized by fulfilling as many of the investor’s goals as possible, weighted by their importance to the investor.

We note that because we use dynamic programming, which evolves backwards in time, the nature of each goal cannot depend on whether other previous goals were realized or not, because that is a forwards in time phenomenon. We see this as the most important limitation of the method in this paper. For example, if an investor has a goal of purchasing a car at a given time, they can either buy the nicest car they considered (full goal) or less nice cars (partial goals) or not buy the car. But should they decide not to buy the car, they cannot then move this car goal to a later year with the method in this paper.

The GBWM algorithm has useful ramifications. First, it optimally prioritizes when to fulfill or forgo any goal, including when it is best to fulfill a partial version of a goal. Instead, forwards-in-time Monte Carlo methods that are widely used in the wealth management industry are restricted to 1) use fixed investment portfolio strategies, such as target date funds or 60/40 strategies, and 2) fulfill as many full, not partial, goals as possible in chronological order. We show that the GBWM algorithm beats traditional methods by achieving goals with higher probabilities, while enabling more flexible financial planning. Second, it better handles mental accounting ([Thaler \(1985\)](#); [Shefrin and Statman \(2000\)](#); [Das et al. \(2010\)](#)), in which investors have different goals in separate mental buckets and express different risk preferences for each mental account. Rather than optimize each mental account separately, the algorithm optimizes all goals in a single dynamic portfolio, enabling offsets between underfunded and overfunded goals. This simplifies and addresses the problem of allocating money across mental accounts (endogenously solved by [Parker \(2021\)](#)) by handling all allocations within a single account. In our framework, wealth does not need to be divided across goals, as is considered in [Parker \(2020\)](#).

Our approach in this paper proceeds as follows: Section 2 explains how the problem is formulated for multiple full and partial goals over a grid of feasible wealth values used at each time period in the algorithm. Section 3 presents the dynamic programming formulation and solution, which determines the optimal investment strategy and optimal goals taking strategy at each wealth grid point and time period. It also shows how to

compute the probability of attaining each goal, including the probability of being at each wealth grid point in each time period, so the investor can fully understand the effect of the optimal strategy on attaining their goals and minimizing their GBWM risk. Section 4 demonstrates a variety of numerical examples from the algorithm. We present both simple examples to give insight into the nature of the solution, as well as complicated examples that are more realistic. We build up to a realistic example (referred to above) in which we use the algorithm to optimize the investing and goals taking strategy for a couple in their mid-thirties over the course of the next 60 years. This couple considers a number of competing annual goals of varying importance that include paying for mortgages, property tax, long-term care insurance, medical expenses, other everyday expenses, cars, house remodeling, trips, philanthropy, and, for their child, orthodontia, private high school tuition, college tuition, and wedding expenses. Finally, in Section 5, we conclude with some final comments and future directions for this work.

2 The Utility (Importance) and Cost of an Investor’s Goals

2.1 Basic Variables and Notation

The basic quantities and variables in the setup are

- **Time:** We consider time periods $t = 0, 1, \dots, T$ with an interval of h years between time periods. So if $h = 0.25$, then $t = 4$ corresponds to one year from the present, $t_0 = 0$. The final time period, $t = T$, for the portfolio may or may not correspond to the projected date of death for the investor.
- **Infusions:** The investor can specify pre-determined wealth infusions, $I(t) > 0$, that they will contribute to their portfolio at any time period or collection of time periods $t = 1, 2, \dots, T - 1$. The values taken by $I(t)$ over these time periods may be chosen to be identical or different. For example, automatic infusions from paychecks may be chosen to remain constant or they may be chosen to be indexed by inflation.
- **Portfolio Evolution and Portfolio Investment Strategies:** The dynamic programming approach works with any Markovian stochastic evolution model for portfolios with deterministic parameters, but for simplicity we will generally use geometric Brownian motion for the evolution model in this paper. (We provide brief results for fat-tailed distributions as well later in the paper.) We assume that the investor has access to l_{\max} different possible portfolio investment strategies, indexed by $l = 1, 2, \dots, l_{\max}$. For the examples in Section 4, we will choose these different portfolios to be along the efficient frontier (see [Markowitz \(1952\)](#)) with the ordering $\mu_1 < \mu_2 < \dots < \mu_{l_{\max}}$ for the portfolios’ expected returns and $\sigma_1 < \sigma_2 < \dots < \sigma_{l_{\max}}$ for the corresponding portfolios’ volatilities. Should the selected portfolios not be on the efficient frontier, they will still conform to the same ordering. A portfolio that cannot fit this ordering should not be used, since it is guaranteed to have both a lower expected return and a higher volatility than at least one of the other

portfolios, and therefore cannot minimize risk from a GBWM (or a traditional) viewpoint. At times we will use the notation μ_{\min} for μ_1 , μ_{\max} for $\mu_{i_{\max}}$, σ_{\min} for σ_1 , and σ_{\max} for $\sigma_{i_{\max}}$.

- **Cost and Utility Vectors:** Implementing full or partial goals at a given time period t results in a reduction in portfolio worth to pay for the goal, with an accretion in utility for the investor. As will be explained in the next two subsections, the potential costs and utilities from implementing combinations of various full and partial goals at a given time period t will be contained in the cost vector $\mathbf{c}(t)$ and the utility vector $\mathbf{u}(t)$, with corresponding components $c_k(t)$ and $u_k(t)$, where $k = 1, 2, \dots, k_{\max}(t)$.
- **Initial Wealth:** We will denote the initial wealth, W , that an investor puts into their portfolio at $t = 0$ by $W(0)$.
- **Wealth Grid:** In Subsection 2.4, we detail the grid of possible wealth values that forms the state space used by the dynamic programming model in Section 3. This wealth grid, which is the same at each time period, contains i_{\max} wealth values, where these wealth values, $W_{\min} = W_1 < W_2 < \dots < W_{i_{\max}} = W_{\max}$, have equal *logarithmic* spacing. Since bankruptcy is a possibility in the examples in this paper, W_{\min} has been chosen to be a small positive number like a dollar or five dollars, which essentially corresponds to being bankrupt.

2.2 Assigning a cost and a utility to each full or partial goal

A classic dynamic programming problem (Merton (1969); Merton (1971)) is to determine how to optimally evolve an investor’s portfolio over time so as to maximize $E[U(W(T))]$, the expected utility of the investor’s wealth, W , at the terminal time $t = T$. A utility function, U , over the continuous domain of wealth values must be selected that corresponds to the investor’s preferences. There are many models for these utility functions as discussed in the introduction. Once a model is assumed, its parameters must also be fit to the individual investor’s preferences. Determining both a model and its parameters is a hard, almost inherently inaccurate process, but it is an unavoidable part of working with utilities when we have a continuum of input values to consider. Our model works with any desired utility function and parameters specified at time T , so we can explore the effect of example models where $U(W(T)) \neq 0$, as we do in Subsection 6.5 of the appendix, but otherwise we set $U(W(T)) = 0$ in this paper, so that we may focus on the effects of the multiple goals that occur at earlier times, $t = 0, 1, \dots, T - 1$. Setting $U(W(T)) = 0$, of course, corresponds to a case where the wealth at the terminal time does not matter, such as when time T corresponds to a projected date of death for an investor who does not care what happens to their money after they die.

At the earlier times, $t = 0, 1, \dots, T - 1$, there is a considerable difference between a goal based perspective, where the investor wishes to purchase specific items, and more standard approaches where the investor looks to maximize a function of annual consumption. For example, a typical standard approach is to optimize $E \left[\sum_{t=1}^T e^{-\beta t} U(C_t) \right]$, where C_t is the yearly consumption, $U(\cdot)$ is a concave, increasing function over the continuum of possible consumption values, and β is the investor’s subjective discount rate. Because

this standard approach involves a continuum of consumption values, we again require a model for the utility function, along with an estimation of both its parameters and the discount parameter β , and again we must accept the fact that it is inherently difficult to match the model and the parameters to any specific investor's preferences.

With multi-goals wealth management, we have a different inherent difficulty, which is that the utility model must accurately reflect the investor's preferences between each of their goals. A single utility function over the continuum of potential goal costs is not appropriate here because two goals can be of different importance to an investor, even if the two goals have the same cost, or, for that matter, a goal that costs less may have more importance to an investor than another goal that is more expensive. Having a different utility function over the continuum of costs for each goal can address this, but it is almost certain not to reflect the investor's actual preferences, given the difficulty of having even one utility function reflect an investor's preferences over a continuum.

Fortunately, none of this is necessary nor desirable. Unlike consumption optimization, multi-goal optimization works with a finite number of goals, as opposed to a continuum of wealth values or consumption values. This means we do not need or want to specify *any* utility function over a continuous domain, meaning we can avoid the considerable assumptions needed to determine the model and parameters for such a function. We instead require appropriately assigning a single utility value to each full or partial goal.

There is no valid way to do this without the investor's feedback. It is inherently central to the process. But can we get *useful* feedback from the investor? Part of the difficulty of determining utility functions over continuous domains is that it is unclear how to map feedback from an investor into a correct utility function. At first that type of a problem appears inherent in our multi-goals formulation as well, since an investor would know that assigning a higher utility to a goal relative to the others will make it more important, but it is unclear how much higher they should choose the assigned utility value to be.

The key to making this work is to note that once we have assigned utility values to each goal, we will be able to use dynamic programming, as described in Section 3, to determine the optimal set of probabilities for attaining each goal given these utility values, and these optimized probabilities are intuitive to investors, meaning they enable the investor to directly give useful feedback. If the resulting optimized probabilities do not fit the investor's preferences, they can boost (or lower) the assigned utilities for their goals, knowing that boosting one goal's utility will increase the probability that it is attained but, in general, decrease the probability that the other goals are attained. This iterative process is necessary to determine the investor's preferences. Any other way would attempt to determine an investor's preferences between goals for them, which is both undesirable and unnecessary. A thorough example of this iterative process will be shown later in Section 4.2.1, but in this subsection we merely need to show what to do with a given set of utilities that are assigned to each goal.

Consider, for example, an investor who has a goal to upgrade to a top-notch electric car four years from now that will cost \$50,000. If h , the time step, is 0.5, then this goal is at $t = 8$ (i.e., in four years) with a cost of $c = \$50,000$. Let's say we assign a utility of $u = 300$ to this goal. This first assigned utility can be any real number. Its value is irrelevant, since it is only the relative value of the goals' utilities to each other that will

matter. As currently stated, this car goal is an example of an all-or-nothing goal, since they either buy or don't by the car. Should the investor be open to alternatives, however, they can consider partial goals. Maybe they are open to a hybrid car with lots of features for \$32,000 or a version with fewer features for \$28,000. Let's say we assign $u = 125$ for the hybrid with lots of features and $u = 80$ for the hybrid with fewer features. Again, there is no need to worry in this subsection if these assigned utility values correctly reflect investor preferences, because they will be changed by the investor later. So, for Goal 1, upgrading the car, we have that

$$\text{Goal 1 } (t = 8): \begin{array}{|c|c|c|c|c|} \hline \text{Cost} & 0 & 28 & 32 & 50 \\ \hline \text{Utility} & 0 & 80 & 125 & 300 \\ \hline \end{array},$$

where the cost is in thousand of dollars. Note that there are four possibilities here: forgoing the goal completely (with a cost and utility of zero), the two partial goals, and the full goal.

The investor may also have a level one goal of paying each semester's tuition for their child at a specific four year college. If tuition currently costs \$30,000 per year, which is projected to increase at a rate of 8% per year, and if the child is intending to start college at $t = 6$, we would have a semi-annual cost of $\$30,000 \cdot 1.08^3 \cdot 2 = \$18,895$ at $t = 6$ and $t = 7$, a cost of $\$30,000 \cdot 1.08^4 \cdot 2 = \$20,407$ at $t = 8$ and $t = 9$, and we then continue in this manner for the college goals at $t = 10, 11, 12$, and 13 . Assume that a utility of 1000 is assigned to each of these eight college goals. We note that this constant utility assignment does not need to be adjusted over time by a discount factor like $e^{-\beta t}$, since its being applied to the inherently subjective worth of tuition to the investor, as opposed to consumption, which is affected by, for example, inflation. At times other than $t = 8$, this tuition goal is called Goal 1, but since we already have the *concurrent* goal of upgrading the car at $t = 8$, the tuition goal at $t = 8$ is labelled to be Goal 2:

$$\text{Goal 2 } (t = 8): \begin{array}{|c|c|c|} \hline \text{Cost} & 0 & 20.407 \\ \hline \text{Utility} & 0 & 1000 \\ \hline \end{array}.$$

When we have concurrent goals, there is some additional processing necessary to remove any illogical goal combinations, as we explain in the next subsection.

2.3 The Cost and Utility Vectors for a Year's Concurrent Goals

Consider an investor who hopes to fulfill three different concurrent goals at a specific time period. Goal 1 is an all-or-nothing goal. Goals 2 and 3 allow for being partially fulfilled. Specifically:

$$\begin{array}{l} \text{Goal 1: } \begin{array}{|c|c|c|} \hline \text{Cost} & 0 & 7 \\ \hline \text{Utility} & 0 & 100 \\ \hline \end{array} \quad \text{Goal 2: } \begin{array}{|c|c|c|c|} \hline \text{Cost} & 0 & 9 & 20 \\ \hline \text{Utility} & 0 & 90 & 300 \\ \hline \end{array} \\ \text{Goal 3: } \begin{array}{|c|c|c|c|c|c|} \hline \text{Cost} & 0 & 10 & 20 & 30 & 40 \\ \hline \text{Utility} & 0 & 40 & 250 & 400 & 500 \\ \hline \end{array}. \end{array}$$

Since there are two possibilities for Goal 1 (fulfill the goal or forgo it), three possibilities for Goal 2, and five possibilities for Goal 3, we have a total of $2 \cdot 3 \cdot 5 = 30$ possibilities for combined goal fulfillment at this time period.

From these 30 possibilities, we first create a table with 30 columns containing the total cost and the total utility for each possibility. These are shown in two rows below, the second row continuing on from the first:

Cost	0	10	20	30	40	9	19	29	39	49	20	30	40	50	60
Utility	0	40	250	400	500	90	130	340	490	590	300	340	550	700	800

7	17	27	37	47	16	26	36	46	56	27	37	47	57	67
100	140	350	500	600	190	230	440	590	690	400	440	650	800	900

We then re-order the table's columns so that the total cost is monotonically increasing. If there are multiple columns corresponding to the same cost amount, we only retain the column that corresponds to the highest utility for that cost amount. This reduces the 30 cases to 24:

Cost	0	7	9	10	16	17	19	20	26	27	29	30	36	37	39
Utility	0	100	90	40	190	140	130	300	230	400	340	400	440	500	490

40	46	47	49	50	56	57	60	67
550	590	650	590	700	690	800	800	900

Finally, starting with the second column, we remove any column where the preceding column has a higher (or equal) utility. We remove these columns because it never makes sense to obtain them, given that the previous column attains a higher (or equal) total utility at a lower cost. This reduces the 24 cases to 13. These 13 cases comprise the final cost and utility vectors:

Cost	0	7	16	20	27	36	37	40	46	47	50	57	67
Utility	0	100	190	300	400	440	500	550	590	650	700	800	900

We note that both the cost vector, $\mathbf{c}(t)$, and the utility vector, $\mathbf{u}(t)$, now contain strictly increasing sequences. As stated in Subsection 2.1, we will use the subscript k to denote the components of these vectors, where $k = 1, 2, \dots, k_{\max}(t)$. So, in the above example, if $k = 3$, then $c_k(t) = 16$ and $u_k(t) = 190$. Also, $c_{k_{\max}(t)}(t) = 67$ and $u_{k_{\max}(t)}(t) = 900$, where $k_{\max}(t) = 13$.

We note that the computer program must retain information for how each of these $k_{\max}(t)$ entries corresponds to the original goals. For example, the computer must retain that the $k = 8$ entry, which has a cost of \$40 and a utility of 550, corresponds to not taking Goal 1, total fulfillment of the Goal 2, and partial fulfillment of the Goal 3 at a cost of \$20.

2.4 The Wealth Grid

Our solution to the multiple goals GBWM problem is implemented at every time period, t , on a grid of wealth values. This grid contains i_{\max} wealth values, where i_{\max} can be any desired number, although suggestions for choosing values of i_{\max} that are not too small to lead to inaccuracies nor too large to unnecessarily slow computations can be found in [Das et al. \(2019\)](#), and we have suitably modified the grid construction approach taken in

that paper for the different types of cashflows here. These i_{\max} wealth values are spread between the wealth limits W_{\min} and W_{\max} .

We begin by approximating W_{\min} and W_{\max} , noting that, ideally, they should represent the lowest and highest possible wealth values reasonably attainable for a solvent investor. Since we have chosen to use geometric Brownian motion for the portfolio evolution model in this paper, we have that an initial wealth, $W(0)$, when affected by no additional external monetary events, will evolve by

$$W(t) = W(0)e^{\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma \sqrt{t}Z},$$

where Z is a standard normal random variable. The effect of any additional external monetary event, be it an infusion, $I(t)$, or the cost of attaining a goal, $c_k(t)$, will also evolve by geometric Brownian motion. We look at the combined effect at each time period of these three types of external monetary events under the best and worst possible scenarios, subject to the assumption that $3 \leq Z \leq 3$. We then locate the lowest and highest computed wealth values from all time periods, which gives us the approximated lower and upper wealth limits:

$$\tilde{W}_{\min} = \min_{t \in \{0, 1, 2, \dots, T\}} \left[W(0)e^{\left(\mu_{\min} - \frac{\sigma_{\max}^2}{2}\right)t + 3\sigma_{\max} \sqrt{t}} + \sum_{\tau=0}^t (I(\tau) - c_{k_{\max}}(\tau)) e^{\left(\mu_{\min} - \frac{\sigma_{\max}^2}{2}\right)(t-\tau) + 3\sigma_{\max} \sqrt{t-\tau}} \right] \quad (1)$$

$$\tilde{W}_{\max} = \max_{t \in \{0, 1, 2, \dots, T\}} \left[W(0)e^{\left(\mu_{\max} - \frac{\sigma_{\min}^2}{2}\right)t + 3\sigma_{\max} \sqrt{t}} + \sum_{\tau=0}^t I(\tau) e^{\left(\mu_{\max} - \frac{\sigma_{\min}^2}{2}\right)(t-\tau) + 3\sigma_{\max} \sqrt{t-\tau}} \right]. \quad (2)$$

Given that most investors will have a number of goals, it will be common for \tilde{W}_{\min} to be negative at this stage. We do not want this, so if $\tilde{W}_{\min} < W_{\text{bankrupt}}$, where W_{bankrupt} is a chosen positive wealth, we set $\tilde{W}_{\min} = W_{\text{bankrupt}}$. Initially, it might seem reasonable to think that W_{bankrupt} should be selected to be one cent, but that can have significant negative implications from a computational point of view. Because we want the points in the wealth grid to have approximately equal *logarithmic* spacing, there will be a concentration of lower wealth values in the grid. This can mean, for example, that choosing W_{bankrupt} to be \$100 instead of one cent can cut the computational time in half, even though having \$100 is essentially the same as being bankrupt from the point of view of an investor whose goals are in terms of thousands, not hundreds, of dollars.

To obtain the logarithmic spacing on the wealth grid, we define \tilde{W}_i by

$$\ln(\tilde{W}_i) = \ln(\tilde{W}_{\min}) + \frac{i}{i_{\max}} \frac{1}{1} \left(\ln(\tilde{W}_{\max}) - \ln(\tilde{W}_{\min}) \right), \text{ where } i = 1, 2, \dots, i_{\max}.$$

We need one of the grid points to equal the initial wealth, $W(0)$, so we shift the entire grid downwards as little as possible to accomplish this. That is, we first define ε to

be the smallest non-negative value with the property that $\ln(\tilde{W}_i) - \varepsilon = \ln(W(0))$ for some value of i , then we define the wealth grid values, W_i , where $i = 1, 2, \dots, i_{\max}$ by $\ln(W_i) = \ln(\tilde{W}_i) - \varepsilon$, or, after exponentiating,

$$W_i = \tilde{W}_i e^{-\varepsilon}.$$

This also generates the wealth limits: $W_{\min} = W_1$ and $W_{\max} = W_{i_{\max}}$.

3 Dynamic Programming for Optimizing in the Case of Multiple Investor Goals

With the setup in place from Section 2 for notation and the cost and utility vectors, we are now prepared to optimize the solution to the GBWM problem with multiple goals by using dynamic programming. Dynamic programming evolves the solution backwards in time, starting from the portfolio's time horizon, $t = T$. For each successive time period at every point on the wealth grid, we determine the strategy for attaining and forgoing goals and for selecting investment portfolios that optimize the total expected utility ultimately collected by the investor. The value function, $V(W_i(t))$, equals this optimized total expected utility, and it is determined by the Bellman (1952) equation, which will be equation (4) in Section 3.2. Once we have used the Bellman equation to determine the optimal strategy at every time and wealth grid point, we then work forwards in time, starting at the initial wealth, $W(0)$, to generate the probability of being at any wealth grid point at any time period. This also generates the probability of fulfilling every full or partial goal.

3.1 Transition Probabilities

We now look to generate the value function at earlier time periods. More specifically, we will use the fact that we know $V(W_i(t+1))$, that is, the value function at all points on the wealth grid at time $t+1$, to determine $V(W_i(t))$, the value function at all points on the wealth grid at time t . This will allow us to iterate backward from time period $T-1$, then to $T-2$, etc., until we end at time 0.

To accomplish this, we need to know the transition probabilities between grid points at time period t to grid points at time period $t+1$. That is, we want to know the probability of transitioning to each value on the wealth grid, $W_j(t+1)$, at time $t+1$ if we are currently at a specific wealth value, $W_i(t)$, at time t and we select a specific $c_k(t)$ from the cost utility vector and a specific investment portfolio corresponding to μ_l and σ_l .

Recall that although we can accommodate any Markovian evolution model with deterministic parameters, for simplicity we have assumed geometric Brownian motion for most of the examples in this paper:

$$W(t) = W(0)e^{\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma Z_t}. \quad (3)$$

Define $\phi(z)$ to be the value of the probability density function of the standard normal random variable at $Z = z$. We adapt equation (3) to these circumstances by noting that we start at time period t with $W_i(t) + I(t) - c_k(t)$ dollars and progress h years to the time period $t + 1$. If we then isolate Z in equation (3), we obtain the approximate transition probabilities, \tilde{q} :

$$\tilde{q}(W_j(t + 1)|W_i(t), c_k(t), \mu_l) = \phi\left(\frac{1}{\sigma_l} \frac{1}{h} \left(\ln\left(\frac{W_j(t + 1)}{W_i(t) + I(t) - c_k(t)}\right) - \left(\mu_l - \frac{\sigma_l^2}{2}\right)h\right)\right).$$

These are approximate because their sum over all the j nodes at time $t + 1$ is not necessarily equal to one. To obtain the transition probabilities, q , which do sum to one, we just normalize the approximate transition probabilities, \tilde{q} :

$$q(W_j(t + 1)|W_i(t), c_k(t), \mu_l) = \frac{\tilde{q}(W_j(t + 1)|W_i(t), c_k(t), \mu_l)}{\sum_{J=1}^{i_{\max}} \tilde{q}(W_J(t + 1)|W_i(t), c_k(t), \mu_l)}.$$

Again, alternative stochastic processes to geometric Brownian motion can be easily accommodated. For example, should fatter tails, which are more realistic, be desired, the specification of Z in equation (3) may be replaced by a random draw from a t-distribution. This is just as easy to calculate. The effect of using a t-distribution will be presented later in Section 4.5. Another possible alternative is to use scenarios generated by matching the moments of the empirical profile of the wealth process. Note that moment matching methods also determine the probability of each scenario, thus making possible the computation of the probability of success. See, for example, Geyer et al. (2014) and Consiglio et al. (2016).

3.2 Bellman Equation

The value function, V , at time $t + 1$ is the highest possible expected sum of the utilities from fulfilled goals in time periods $t + 1, t + 2, \dots, T - 1$ plus the utility of the final excess wealth at $t = T$. The term “highest possible” means that at each time and point on the wealth grid, the optimal choices for the component k from the cost/utility vectors and for the investment portfolio l are used.

At time T , we know from our discussion in Subsection 2.2 that $V(W_i(T)) = U(W_i(T))$. Since we know the transition probabilities from Subsection 3.1, we can compute $V(W_i(T - 1))$ by substituting $t = T - 1$ into the following Bellman equation and solving it for each $i = 1, 2, \dots, i_{\max}$:

$$V(W_i(t)) = \max_{k,l} \left[u_k(t) + \sum_{J=1}^{i_{\max}} V(W_J(t + 1)) q(W_J(t + 1)|W_i(t), c_k(t), \mu_l) \right]. \quad (4)$$

The values of k and l that correspond to the maximum in the Bellman equation are the optimal choices, which lead to the highest possible expected utility sum. We can then repeat this process for the time period $t = T - 2$, then $t = T - 3$, etc., stopping after the results at $t = 0$ are computed. This gives the value function, and more importantly the optimal strategy, at all times and all gridpoints. As we progress through this Bellman

procedure, at each gridpoint, i , and time, t , we store these optimal strategy choices for k and l , labeling them $k_{i,t}$ and $l_{i,t}$.

The overall runtime increases as $k_{\max}(t)$, l_{\max} , the number of grid points i_{\max} , and the number of periods T increase. We note that in the context of this paper, when $k_{\max}(t)$ and l_{\max} are both large, optimizing the Bellman equation over $k_{\max}(t) \times l_{\max}$ possibilities slows down the algorithm, in which case it is computationally wiser to split each time period optimization into two parts: first over the l_{\max} portfolios, then over the $k_{\max}(t)$ goal choices. This essentially leads to optimizing over $k_{\max}(t) \times l_{\max}$ possibilities, instead of $k_{\max}(t) \times l_{\max}$ possibilities. For each portfolio we have to compute transition probabilities for all wealth values on the grid, which are i_{\max} in number. Since we do this in a vectorized fashion, it is optimized — we use NumPy vectors in the Python programming language, further optimized using just-in-time (jit) compilation — so we may assume no additional complexity from this step. Therefore, the complexity at each time t is of order $O(i_{\max}(k_{\max}(t) + l_{\max}))$. Given there are T periods in the model, we end up with an overall complexity of $O\left(i_{\max}\left(T l_{\max} + \sum_{t=0}^T k_{\max}(t)\right)\right)$.

For all but the experiment in Subsection 4.5, the computational time was under 2 seconds, so the split optimization approach was unnecessary. However, the example in Subsection 4.5 has $l_{\max} = 15$ possible portfolios, $i_{\max} = 1221$ wealth grid points, and 301 goals and an additional 138 partial goals spread out over $T = 60$ years. Without the split optimization approach, the computational time for this example was about 13 minutes; with the split optimization approach, the computational time was reduced to just 17 seconds on a regular desktop computer.

3.3 Probability Distribution for Wealth and Optimally Attaining Goals

To determine the probability distribution for the investor’s wealth at future times, we use the transition probabilities and the optimal strategy information, $k_{i,t}$ and $l_{i,t}$, determined from the Bellman equation to evolve the probability distribution forward in time, starting with $t = 0$, then $t = 1$, and ending with $t = T - 1$.

More specifically, at $t = 0$, define i_0 so that W_{i_0} is the wealth node that equals $W(0)$, therefore $p(W_{i_0}(0)) = 1$ and $p(W_i(0)) = 0$ for all $i \neq i_0$. To obtain the wealth probability distribution at $t = 1$, we set $t = 0$ in the following “forwards equation” and run it for each $j = 1, 2, \dots, i_{\max}$:

$$p(W_j(t+1)) = \sum_{i=1}^{i_{\max}} q(W_j(t+1)|W_i(t), c_{k_{i,t}}(t), \mu_{l_{i,t}}) p(W_i(t)). \quad (5)$$

We then set $t = 1$ in equation (5) and again run it for each $j = 1, 2, \dots, i_{\max}$, continuing in this manner until we finish with the $t = T - 1$ case. This gives the probability distribution for every point in the wealth grid at every time period of the portfolio.

Once this wealth probability distribution is calculated, we can determine the probability of attaining any specific full or partial goal at any given time t : At each time t for

each $k = 1, 2, \dots, k_{\max}(t)$, we sum $p(W_i(t))$ over every i where $k_{i,t} = k$. This gives the probability that each component k in the cost/utility vectors will be chosen. Once this is known, the cost/utility vectors are reconnected to their original goals, and the probability that each full or partial goal will be fulfilled is determined by summing over the components connected to that full or partial goal. For example, if the cost/utility vector components $k = 2, 5$, and 9 are the only entries corresponding to, say, total fulfillment of goal two, and their probabilities are 0.05 , 0.07 , and 0.04 , then the probability of totally fulfilling goal two is $0.05 + 0.07 + 0.04 = 0.16$.

This goal probability information can then be given to the investor. If the investor finds these goal probabilities do not fit well with their preferences, they may change the utility values assigned to their goals or increase their infusions and then run the dynamic programming algorithm again to see if the new results sufficiently meet their desires. (See subsection 4.2.1 for an example of this.) This iterative process enables the investor to gain a thorough understanding of what the trade-offs are among their goals when using optimal goal fulfillment and optimal portfolio investment strategies.

4 Numerical Examples

In this section, we present examples that demonstrate the features of the approach, including how changing inputs affects the optimal strategy and other results determined by the algorithm. We start with simple cases where the features are clearer and then proceed to more complex, realistic situations.

In the examples, the portfolios that are available to the investor will be on the efficient frontier shown in Figure 1. This frontier was generated from historical returns in the 20 year period between January 1998 to December 2017 for index funds representing U.S. Bonds, International Stocks, and U.S. Stocks¹.


Further, unless stated otherwise, for the examples in this section we will assume that

- We restrict the portfolios used on the frontier to be between $\mu_{\min} = 0.0526$ (which corresponds to $\sigma_{\min} = 0.0374$) and $\mu_{\max} = 0.0886$ (which corresponds to $\sigma_{\max} = 0.1954$). The point $(\sigma_{\min}, \mu_{\min})$ corresponds to the vertex on the frontier. The value of μ_{\max} corresponds to the highest of the three component index fund returns, specifically, the return on U.S. Stocks. These three asset classes were chosen because they are widely used by wealth management firms in constructing target date funds.
- We set the number of available investment portfolios, $l_{\max} = 15$, and these 15 portfolios on the frontier have μ values that are equally spaced between μ_{\min} and μ_{\max} .

¹The three index funds used are (i) Vanguard Total Bond Market II Index Fund Investor Shares (VTBIX), representative of U.S. Fixed Income (Intermediate-Term Bond), (ii) Vanguard Total International Stock Index Fund Investor Shares (VGTIX), representative of Foreign Equity (Large Cap Blend), (iii) Vanguard Total Stock Market Index Fund Investor Shares (VTSMX), representative of U.S. Equity (Large Cap Blend). These three funds have been chosen as representatives of their respective asset categories for illustrative purposes only.

- The time step $h = 1$, so the value of t corresponds to the number of years after the initial investment at $t = 0$ is deposited.
- We assume no bequest motive, so $V(W_i(T)) = U(W_i) = 0$ for all wealth grid points W_i . However, this is without loss of generality and we provide examples with bequest utility in Section 6.5.
- There are no infusions, which means $I(t) = 0$ at all times t .

These assumptions may be easily changed if desired, as we will do in some examples later in this section.



Portfolio number	Portfolio Weights		
	U.S. Bonds	International Stocks	U.S. Stocks
1	0.9098	0.0225	0.0677
2	0.8500	0.0033	0.1467
3	0.7903	-0.0160	0.2257
4	0.7305	-0.0352	0.3047
5	0.6707	-0.0545	0.3837
6	0.6110	-0.0737	0.4628
7	0.5512	-0.0930	0.5418
8	0.4915	-0.1122	0.6208
9	0.4317	-0.1315	0.6998
10	0.3719	-0.1507	0.7788
11	0.3122	-0.1700	0.8578
12	0.2524	-0.1892	0.9368
13	0.1927	-0.2085	1.0158
14	0.1329	-0.2277	1.0948
15	0.0731	-0.2470	1.1738

Figure 1: Top: The efficient frontier generated from the returns of the three index funds. Bottom: The portfolio weights in the three index funds for each of the $l_{\max} = 15$ portfolios with equal spacing in μ between $\mu_{\min} = 0.0526$ and $\mu_{\max} = 0.0886$. As seen from the table, both long-only and long-short portfolios occur.

Our dynamic programming algorithm is easily implemented on a single computer (desktop or laptop) and is coded up in the Python programming language. With the machine we used, which has 16 GB of RAM and an i5 Intel CPU, we will see that the

runtime stays under 2 seconds, except for the complicated final example in Subsection 4.5, where the runtime is 17 seconds.

This section contains several subsections with many examples in a variety of settings and scenarios. We therefore provide a roadmap for these examples in Table 1.

4.1 Two All-or-Nothing Goals

For the examples in this section, we have just two goals: one at $t = 5$ years, the other at $t = 10$. (We therefore use $T = 11$, with no bequest at T .) We use an initial wealth of $W(0) = \$100$, which leads to $W_{\max} = \$1834$. We have selected $i_{\max} = 475$ for the number of nodes in the wealth grid between $W_{\min} = 1$ and W_{\max} . The runtime for this model is 0.5 seconds.

4.1.1 A Single Case with Two All-Or-Nothing Goals

We first consider a case where the goal at $t = 5$, if fulfilled, has a cost $c(5) = \$100$ and an assigned utility $u(5) = 1000$, while the goal at $t = 10$ has a cost $c(10) = \$150$ and utility $u(10) = 1000$. Because the utilities of the two goals are equal, it is clear what the investor should do regarding fulfilling the goal at $t = 5$: If the investor has the \$100 at $t = 5$, they should spend it to attain the $t = 5$ goal, because this guarantees a utility of 1000 with the possibility for an additional 1000 should they be able to obtain \$150 when $t = 10$. In contrast, had they chosen to forgo the goal at $t = 5$, they would, at most, have a total utility of 1000 from the $t = 10$ goal, which corresponds to a smaller total expected utility.

It is far less clear for this case which investment portfolio choices optimize the investor’s total expected utility, but the algorithm generates these optimal choices, which are given in Figure 2, for all levels of wealth (y -axis) and values of time (x -axis). In this figure, the darker the portfolio color at a given wealth and time, the higher the optimal portfolio is located on the efficient frontier, meaning the portfolio is more aggressive.

We see interesting strategy shifts in Figure 2 when $t < 5$. If the portfolio is doing poorly, the most aggressive portfolio is chosen, since that optimizes the chance of the investor getting back to \$100 and attaining the $t = 5$ goal. If the investor has a little more money, they optimally choose relatively safer portfolios to make sure they remain on-track to have at least \$100 at $t = 5$. But if they have even more than that, they optimally shift to more aggressive portfolios again since they are likely to attain the $t = 5$ goal and now hope to also attain the $t = 10$ goal. Finally, should they have even more money, they again move to safer portfolios, since they now want to preserve the wealth that will allow them to attain both goals.

At $t = 5$, the most aggressive portfolio is chosen for wealth levels below approximately \$201 because \$100 will be deducted if the investor has that much money, leaving the need for a more aggressive portfolio to optimize the chance of growing the remaining \$101 or less into \$150 by $t = 10$.² For $t > 5$, unsurprisingly, we see aggressive portfolios if

²At the wealth level barely above \$100, we see a light colored pixel. In this case, the \$100 goal is met,

Table 1: Roadmap of examples using the multiple goals algorithm.

Experiment	Section	Table/Figure	Implementation Plan	Intent & Outcome
Two all or nothing goals (10 years, 2 goals)	Sec 4.1.1	Table 2, Figures 2, 3	Explain optimal investment portfolio choices for various (t, W) combinations; look at wealth distributions at both goals' times	Portfolio choice is sensitive to wealth level around the goal amount; wealth distribution clusters probability mass just above needed goal costs
Varying utilities for two all or nothing goals (10 years, 2 goals)	Sec 4.1.2	Table 2	See how goals are traded off based on costs and utilities	Goal probabilities are sensitive to utilities in intuitive ways
Multiple all or nothing goals (24 years, 7 goals)	Sec 4.2, 4.2.1; Appendix: 6.1, 6.2, 6.3, 6.4, 6.5	Tables 3, 4; Appendix: Tables 10, 11, 12, 13, Figures 5, 6	Examine how multiple goals tradeoff; how changing utilities affects probabilities; iteratively assign utilities. Appendix: comparative statics from changing one goal's utility; effect of utility from terminal wealth	Increasing utilities for one goal increases its probability but it does not mean all other goal probabilities fall; an iterative example of how the investor can arrive at optimal goal utilities. Appendix: Intuition regarding changing assigned utilities to goals, initial wealth, infusions, available investment portfolios, and weight given to terminal wealth utility
Concurrent goals and partial goals (10 years, 3 goals)	Sec 4.3.1	Table 5	Implement the concurrent and partial goal example in Sec 2.3	We see how the tradeoff between partial and full realization of a goal occurs; show how probabilities for both are easily reported
Concurrent goals and partial goals (7 years, semi-annual, 2 goals)	Sec 4.3.2	Table 6	Implement the concurrent and partial goal examples in Sec 2.2; run an example with half-year subperiods	We see how the tradeoff between partial and full realization of a goal occurs; show how probabilities for both are easily reported
Comparison of this paper's algorithm with Target Date Funds (TDF) (60 years, 3 goals)	Sec 4.4	Tables 7, 8	Compare the performance of this paper's algorithm with TDF	This paper's algorithm significantly outperforms TDF because it is time, wealth, and goal optimized, whereas TDF is only time optimized
Long-range financial planning (60 years, hundreds of goals)	Sec 4.5	Table 9	Does this paper's algorithm scale to large, realistic financial planning problems? How do goal probabilities change with goal tier?	Good scaling outcomes with low latency; Goal probabilities are highest for Tier 1, then Tier 2, etc., as desired
Comparison of good and bad realizations	Sec 4.5.1	Figure 4	Examine how different the utilities and goal realizations are for a good portfolio path versus a poor portfolio path	Difference in utilities is small because the algorithm forgoes minor goals when the path is poor




Figure 2: The grid of optimal portfolios at all levels of wealth (y -axis) and time (x -axis). The darker the portfolio color, the higher up on the efficient frontier the portfolio is, meaning both a higher expected return and a higher volatility. As shown in the color bar on the right of the plot, we consider a total of 15 possible portfolios along the efficient frontier, numbered from 0, the lightest color and therefore most conservative portfolio, to 14, the darkest color and therefore the most aggressive portfolio. The cost of the goals are \$100 at 5 years and \$150 at 10 years. Both goals have an assigned utility of 1000.

but there is so little wealth remaining that it is below the lowest value in our logarithmically spaced wealth grid. This state of essential bankruptcy means that the $t = 10$ goal will not be met no matter

reaching \$150 by $t = 10$ looks difficult and less aggressive portfolios when the investor has more money and is better off safeguarding their funds from losses that could reduce their wealth below \$150 by $t = 10$.

This optimal strategy leads to a probability of 0.893 for fulfilling the $t = 5$ goal and a 0.275 probability for fulfilling the $t = 10$ goal. The probability distributions for the wealth at each time are given in Figure 3. Note the figure's leftward shifts after $t = 5$ and $t = 10$ due to the payments made when fulfilling goals. Also, as expected, we see the optimized strategy create a clump in the distribution that just exceeds \$100 at $t = 5$ so as to attain the $t = 5$ goal and another clump exceed \$150 at $t = 10$. Because the probability of exercising the goal at 5 years is much higher than the one at 10 years, we see a bigger probability mass clump at 5 years. Further, note the probability mass clump located a little below \$250 at year $t = 5$, which occurs because the algorithm may be on track to attain both goals.




Figure 3: Probability distribution for wealth at different times, t , in the model. The plot is divided into two time intervals: the first 5 years are contained in the top panel and the remaining years in the bottom panel. Attaining the goals at either $t = 5$ or $t = 10$ causes the distribution to shift to the left. Note how the optimized investing strategy makes the $t = 5$ distribution clump just above $W = 100$ to attain the $t = 5$ goal and then makes the $t = 10$ distribution clump just above $W = 150$ to attain the $t = 10$ goal.

4.1.2 Altering the Two Goals' Assigned Utilities and Costs

The utility values assigned by the investor to the two goals have a significant effect on the optimal portfolio strategy. They also have a significant effect on deciding whether or not to exercise the option to fulfill or forgo the $t = 5$ goal. This decision is optimally determined by the dynamic programming algorithm by calculating which is larger: 1)

which investment portfolio is chosen. In these essentially bankrupt cases, the program defaults to the lightest color.

the utility for the $t = 5$ goal plus the expected utility of the $t = 10$ goal starting from a wealth level that has been reduced due to paying for the $t = 5$ goal or 2) the utility of the $t = 10$ goal with no reduction in wealth because the $t = 5$ goal is not exercised.

In Panel A of Table 2, we see some of the changes and tradeoffs that occur as we vary the utilities of the two goals while keeping the costs fixed at $c(5) = \$100$ and $c(10) = \$150$. In Panel B of Table 2, we look at the same utility combinations, but we switch the costs so that $c(5) = \$150$ and $c(10) = \$100$. An analysis of the table reveals a number of properties concerning the optimal strategy:

Table 2: The tradeoffs between a goal at $t = 5$ and a goal at $t = 10$, each of which is fulfilled or forgone. In Panel A, fulfilling the $t = 5$ goal costs \$100 and fulfilling the $t = 10$ goal costs \$150. In Panel B, these goal costs are switched. For various values of the utility, u , assigned to the two goals, we present the optimal expected total utility, $E[u]$, which is the value function at the initial wealth and time, $W_{i_0}(0) = \$100$. Under this optimal strategy, we present the probability of having sufficient wealth to attain the goal at $t = 5$, that is, $P(W(5) \geq \$100)$ for Panel A and $P(W(5) \geq \$150)$ for Panel B, and for both panels we present the probabilities that the optimal strategy will fulfill the goal at $t = 5$ and $t = 10$, denoted $P(\text{fulfill } t = 5 \text{ goal})$ and $P(\text{fulfill } t = 10 \text{ goal})$.

Panel A: $c(5)=\$100, c(10)=\150

$u(5)$	$u(10)$	$E[u] = V(W_{i_0}(0))$	$P(W(5) \geq \$100)$	$P(\text{fulfill } t = 5 \text{ goal})$	$P(\text{fulfill } t = 10 \text{ goal})$
1000	1000	1168	0.893	0.893	0.275
1000	2000	1855	0.917	0.123	0.866
1000	3000	2757	0.969	0.009	0.916
2000	1000	2110	0.969	0.969	0.171
2000	2000	2336	0.893	0.893	0.275
2000	3000	2886	0.849	0.298	0.764
3000	1000	3087	0.984	0.984	0.134
3000	2000	3259	0.950	0.950	0.205
3000	3000	3504	0.893	0.893	0.275

Panel B: $c(5)=\$150, c(10)=\100

$u(5)$	$u(10)$	$E[u] = V(W_{i_0}(0))$	$P(W(5) \geq \$150)$	$P(\text{fulfill } t = 5 \text{ goal})$	$P(\text{fulfill } t = 10 \text{ goal})$
1000	1000	1185	0.398	0.398	0.787
1000	2000	2137	0.358	0.186	0.976
1000	3000	3120	0.340	0.163	0.986
2000	1000	1631	0.493	0.493	0.644
2000	2000	2370	0.398	0.398	0.787
2000	3000	3306	0.373	0.210	0.962
3000	1000	2155	0.544	0.544	0.524
3000	2000	2792	0.449	0.449	0.723
3000	3000	3555	0.398	0.398	0.787

- As we would expect, the relative utilities for the two goals determines their priority. Holding $u(5)$ constant and then increasing $u(10)$ from 1000 to 2000 to 3000, we see from the last two columns in the table that the probability of attaining the $t = 5$ goal decreases, while the probability for the $t = 10$ goal increases. This can be drastic. For example, in the case where $u(5) = 1000$ in Panel A, we see the probability of

attaining the $t = 5$ goal decrease from 89.3% to 0.9% while the percentage for the $t = 10$ goal increases from 27.5% to 91.6%.

- Changing the cost has a similar effect. Comparing Panel A to Panel B, we see a considerable decrease in the probability of attaining the $t = 5$ goal and a corresponding increase in the probability of attaining the $t = 10$ goal caused by $c(5)$ increasing from \$100 to \$150 and $c(10)$ decreasing from \$150 to \$100.
- The third column gives the value function, V , at the initial condition, $W_{i_0}(0)$, which is equal to the optimal total expected utility. This value can be computed from the first two and last two columns; for example, for the second row in panel A, we can calculate $1855 = 1000 \cdot 0.123 + 2000 \cdot 0.866$. That is, the total expected utility is the sum of each goal's utility weighted by the probability of attaining that goal.
- For any set of goals, multiplying all the goals' utilities by the same constant has no effect on the optimal strategy. So, for the cases with two goals here, only the ratio of $u(5)$ to $u(10)$ matters, not the individual values of $u(5)$ and $u(10)$. For example, within each panel, when $u(5) = u(10)$ we see that the probabilities in the last three columns are identical.
- If $u(5) \geq u(10)$, then the investor should always take the $t = 5$ goal if they have the money. This is due to the fact that taking the $t = 5$ goal guarantees an amount greater than or equal to the utility that may or may not happen from the $t = 10$ goal. This means the values in the fourth and fifth columns of the table are identical. When $u(5) < u(10)$, however, it may be best to forgo the $t = 5$ goal to optimize the chance of fulfilling the $t = 10$ goal. For example, in Panel A when $u(5) = 1000$ and $u(10) = 3000$, there is a 96.9% chance that the investor will have retained at least their original investment of \$100 by $t = 5$, but they optimally opt to spend this \$100 to fulfill the $t = 5$ goal only 0.9% of the time, because it is far more important to build up at least \$150 by $t = 10$ where they attain three times more utility. They only opt to fulfill the $t = 5$ goal if they have amassed so much money that they are confident they can obtain the $t = 10$ goal even after spending \$100 to obtain the $t = 5$ goal.

4.1.3 Determining the Wealth Intervals where the Investor Optimally Chooses to Fulfill the $t = 5$ Goal

Recall from the introduction the case of an investor that wanted to take a nice vacation in 5 years and also wanted to buy their child a car in 10 years, where this second goal was more important and more expensive than the first goal. This corresponds to a situation like $c(5) = 100$, $c(10) = 150$, $u(5) = 2000$, and $u(10) = 3000$, given in Panel A of Table 2. In the introduction we pointed out that at $t = 5$, the investor should repeatedly shift their decision about whether or not to take the goal at $t = 5$ depending on their wealth, but it was not clear at which wealth values these decision switches should occur.

Our algorithm can determine these wealth values for switching the decision optimally. For example, take the case above where $c(5) = 100$, $c(10) = 150$, $u(5) = 2000$, and $u(10) = 3000$. At $t = 5$, if the investor has less than \$100, obviously they cannot take the

$t = 5$ goal. If they have between \$100 and \$108 at $t = 5$, however, the investor should take the $t = 5$ goal, because their wealth is so low that the chance of being able to attain \$150 by $t = 10$ is too small to justify the additional utility they would obtain should they reach the $t = 10$ goal. Between \$108 and \$182, the situation flips: the investor should forgo the $t = 5$ goal, because it sufficiently increases the likelihood of attaining the $t = 10$ goal to justify the risk. Finally, if the investor has over \$182 at $t = 5$, they should again take the $t = 5$ goal, because they are sufficiently likely to be able to attain both goals.

4.2 Multiple All-or-Nothing Goals

In this section, we expand from two all-or-nothing goals to seven all-or-nothing goals whose details are listed in Table 3 below, where $T = 25$. We start with an initial wealth of $W(0) = \$30$ (thousand), which corresponds to $W_{\max} = \$5026$, just over five million dollars. We have selected $i_{\max} = 556$ for the number of nodes in the wealth grid between $W_{\min} = 1$ and W_{\max} . The runtime is approximately 1.5 seconds.

Table 3: Base Case: Seven All-Or-Nothing Goals Under the Optimal Strategy where $W(0) = 30$.

	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
Goal time, t (in years)	5	8	10	11	17	22	24
Goal cost, c (in 1000s of \$)	25	17	15	80	50	60	130
Goal utility, u	1000	2500	500	1500	300	3000	2000
Probability of fulfilling goal	0.0446	0.9871	0.2077	0.0316	0.0192	0.7569	0.2373

Table 3 also provides the results from the algorithm, which optimizes the total expected utility. The value of this optimal total expected utility $E[u] = V(W_{i_0}(0))$ is 5415. As indicated in the previous section, this equals the dot product of the vectors comprising the bottom two rows of Table 3. Since the total utility from summing the elements in the third row is 10800, we see that the investor can adopt a strategy to attain goals that, at best, on average, achieves the fraction $5415/10800 = 0.5014$ of the utility corresponding to attaining all the investor’s goals. In the remainder of this paper, we will generally refer to this optimally achievable total expected utility fraction as the “expected utility fraction.”

In the appendix (Section 6), we explore comparative statics for this example, namely, the effect when the investor considers changing utility assignments, changing the initial wealth, making periodic cash infusions, increasing the range along the efficient frontier for the available portfolio choices, and adopting a non-zero bequest utility. In the next subsection, we explore how an investor might combine the first three of these ideas — making successive changes to their utility values, as well as to their initial contribution and to their later infusions — so as to achieve their desired probabilities of attaining the goals for this example.

4.2.1 The Process of Assigning Utilities to Goals

The GBWM algorithm rigorously enables a “human-in-the-loop” mechanism by which investors can intuitively tweak their utilities and contributions to determine the smallest amount of resources needed to achieve minimum desired probabilities for attaining each of their goals. Determining these minimum probabilities, or potentially altering them, is, by necessity, something the investor must determine. Similarly, changing available resources, like increasing the initial investment or committing to later cash infusions, is also a necessarily human decision. However, because the dynamic programming algorithm in Sections 2 and 3 rigorously translates a given set of goal utilities and financial resources into an optimal set of probabilities for attaining each goal, and these probabilities, unlike the underlying utility values, are intuitive to the investor, they can knowledgeably make alterations to conform to their desires, as we exemplify in this subsection. The iterative, experimental aspect of this process is essential to making certain the investor controls the algorithm to best meet their desires.

We consider a situation where Goals 2 and 6 of the seven goals are the most important to the investor, so they wish to maintain at least a 99% chance of attaining them; Goals 1, 4, and 7 are next most important, where they want to maintain at least a 60% chance; and finally, Goals 3 and 5 are least important, where they want to maintain at least a 30% chance. The optimal probabilities in Table 3, which are restated in Run #1 of Table 4, show that none of these desired probabilities have been achieved. That is, we know that even with the optimal strategy, the investor is unable to attain their desired probabilities, so they must make one of the three changes mentioned above: increase the initial contribution, increase infusions, or expand the possible portfolios available. Assume the investor determines that they can initially invest another \$20 (thousand), so $W(0)$ now becomes \$50 (thousand). The effect of this larger initial investment on the increased probabilities for achieving the seven goals is shown in Run #2 in Table 4.

However, even in Run #2, only two of the seven goals have sufficient probabilities, indicating the investor will likely need to make further changes, such as committing to infusions over time. In preparation for this, the investor alters their assigned utility values by decreasing the utility values for the two goals with sufficient probability and increasing the utility for the goals whose probability is significantly below the desired values, generating Run #3 in Table 4. Since the optimal probabilities for all the goals are again below the desired values, we consider the effect of infusions.

The investor begins by considering an annual infusion of $I(t) = \$2$ (thousand) dollars, which leads to the values seen in Run #4. These are close to the desired probability values. By tweaking the utility values, the investor obtains Run #5, but, once again, sees that the optimal probabilities are all below the desired values, so a small amount of additional funds are still required. The investor therefore decides to increase their annual infusion to \$3 (thousand) dollars in each of the first six years, while keeping the annual infusion at \$2 (thousand) dollars after that. This gives the results in Run #6, and a small tweak of the utility values after that achieves all the desired probabilities, as shown in Run #7.

This example is typical of how an investor, or a robo-advisor program, can use the algorithm to determine the final assigned utility values for each goal. That is, the investor

or program increases the utility of a goal whose probability is too low and decreases the utility of a goal whose probability is too high, until all the probabilities are too low or all the probabilities are high enough. If all the probabilities are too low, the investor must either decrease their expectations by allowing lower probabilities for at least some of their goals, or they must increase the initial investment, infusions, or available portfolios. If all the probabilities are high enough, the investor knows they have devoted sufficient resources to attain their goals, as long as they follow the optimal strategy determined in Sections 2 and 3. This process prevents investors from having unrealistic expectations for what their money can accomplish, and it also prevents the opposite mistake of devoting too many resources to attain their goals.

Table 4: Determining appropriate utility values for the seven goals in Table 3, along with appropriate values for the initial investment and infusions in each future year. Probability values that are sufficiently high for the investor are in boldface.

Minimum goal probability	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
Run #1: Goal utility, u	1000	2500	500	1500	300	3000	2000
Probability of fulfilling goal	0.0446	0.9871	0.2077	0.0316	0.0192	0.7569	0.2373
Run #2: Goal utility, u	1000	2500	500	1500	300	3000	2000
Probability of fulfilling goal	0.3380	0.9994	0.5299	0.1586	0.0761	0.9352	0.4772
Run #3: Goal utility, u	1200	1700	400	2500	700	3000	2000
Probability of fulfilling goal	0.2980	0.9729	0.1321	0.3073	0.2070	0.9338	0.4255
Run #4: Goal utility, u	1200	1700	400	2500	700	3000	2000
Probability of fulfilling goal	0.6832	1.0000	0.2344	0.4982	0.2886	0.9901	0.5803
Run #5: Goal utility, u	1200	800	425	2650	705	2900	2100
Probability of fulfilling goal	0.5645	0.9896	0.2618	0.5478	0.2871	0.9895	0.5881
Run #6: Goal utility, u	1200	800	425	2650	705	2900	2100
Probability of fulfilling goal	0.6476	0.9942	0.3027	0.5862	0.3126	0.9912	0.6106
Run #7: Goal utility, u	1200	760	425	2700	700	2900	2070
Probability of fulfilling goal	0.6237	0.9914	0.3150	0.6046	0.3072	0.9902	0.6027

4.3 Concurrent Goals and Partial Goals

Sections 4.1 and 4.2 were concerned with all-or-nothing goals from different years. We now consider what happens when we allow for multiple goals to happen concurrently and when we allow for partial fulfillment of goals.

4.3.1 Results for the Concurrent and Partial Goals Example in Subsection 2.3

We revisit the example with three concurrent goals described in Subsection 2.3. We will assume these three goals occur at year 5. To this, we add another goal in year 10 that may be taken partially at a cost of \$50 with a utility of 500 or taken fully at a cost of \$90 with a utility of 1000. We assume the initial wealth, $W(0)$, is \$50, so $T = 11$ and

$W_{\max} = \$909$. We have selected $i_{\max} = 419$ for the number of nodes in the wealth grid between $W_{\min} = 1$ and W_{\max} . The runtime is approximately 0.5 seconds.

The optimal total utility is 1013 out of a maximum possible of 1900, so the expected utility fraction is $1013/1900 = 0.5330$. Table 5 shows the probability of reaching various partial and full goals. The final three columns in Panel A show, respectively, the cost vector, $\mathbf{c}(5)$; the utility vector, $\mathbf{u}(5)$; and the probability of fulfilling the partial or full goals associated with the components of these vectors. The first six columns break down the cost and utility vectors' values by goal. This is key to creating Panel B, which gives the probabilities of attaining each full or partial goal as an investor would want to see. For example, the first row in Panel B gives the probability of not attaining Goal 1 in year 5. This is calculated by summing all the probabilities in Panel A that correspond to Goal 1 not being attained: $0.0198 + 0.1547 + 0.0672 + 0.1162 = 0.3579$.

The optimal decisions shown in Panel B make intuitive sense. At year $t = 5$, the investor is first attracted to fully realizing Goal 2, because its utility per cost ratio is very high. The investor may defer this goal so as to optimize the chance of attaining the very high utility values assigned to Goal 1 at $t = 10$, even though these high utility values also correspond to higher costs. If the cost of \$20 for fully attaining Goal 2 is too much, the investor may decide to partially fulfill Goal 2 and totally fill Goal 1 at $t = 5$ for a total cost of \$16. Should the investor have even less portfolio wealth and needs to choose between partially filling Goal 2 or totally filling Goal 1, it is better to choose the latter because it gives a higher utility (100 versus 90) at a lower cost (\$7 versus \$9). Should the investor have enough money at $t = 5$ to fully attain Goal 2 and spending more money at $t = 5$ will not sufficiently hurt the chances of attaining Goal 1 at $t = 10$, the algorithm will look to also fulfill Goal 1 at $t = 5$ and then progressively fulfill more and more of Goal 3 at $t = 5$.

4.3.2 Results for the Example in Subsection 2.2 with a New Car Goal and a College Tuition Goal

We revisit the example from Subsection 2.2 where an investor is considering buying a new car in four years, but also places a much higher utility on paying their child's tuition every semester for four years, starting in year three. Because the tuition is paid semi-annually, we use $h = 0.5$ instead of $h = 1$, so, for example, year 4 corresponds to $t = 8$. Using the method detailed in Subsection 2.3, we can determine the cost vectors, $\mathbf{c}(t)$, and the utility vectors, $\mathbf{u}(t)$, from the numbers given in Subsection 2.2. For year 4 (i.e., $t = 8$), these vectors, $\mathbf{c}(8)$ and $\mathbf{u}(8)$, are given, respectively, in the fifth and sixth columns of Table 6.

We assume the initial wealth, $W(0)$, is \$100, so $T = 14$ (for a horizon of 7 years) and $W_{\max} = \$868$. We have selected $i_{\max} = 594$ for the number of nodes in the wealth grid between $W_{\min} = 1$ and W_{\max} . The runtime is under 1 second. The optimal total utility is 6356 out of a maximum possible of 8300, so the expected utility fraction is $6356/8300 = 0.7658$. Table 6 shows the probability of reaching the various partial and full goals. From Panel B, we see that the tuition goal in every semester is fulfilled with a reasonable probability ($> 75\%$, except in the last year). However, it is almost certain that we will forgo the car goal in year 4, because it is costly and does not offer a utility

Table 5: Partial and concurrent goals. We implement the three goals described in Subsection 2.3 as goals at year 5 and add to that a single goal in year 10 that may be taken partially at a cost of \$50 with a utility of 500 or fully at a cost of \$90 with a utility of 1000. The initial wealth, $W(0)$, is \$50. The detailed breakout of goals and probabilities at year 5 is shown in Panel A. Note that the seventh and eighth columns are, respectively, the cost vector, $\mathbf{c}(5)$, and the utility vector, $\mathbf{u}(5)$. In Panel B we show the costs and utilities for each partial or full goal at each time, along with the probability that the goal is fulfilled. The information in Panel A is summed to obtain the year 5 information shown in Panel B.

<i>Panel A: Year 5</i>								
Goal 1	Cost		Utility			Total		
	Goal 2	Goal 3	Goal 1	Goal 2	Goal 3	Cost	Utility	Probability
0	0	0	0	0	0	0	0	0.0198
7	0	0	100	0	0	7	100	0.0554
7	9	0	100	90	0	16	190	0.0012
0	20	0	0	300	0	20	300	0.1547
7	20	0	100	300	0	27	400	0.2240
7	9	20	100	90	250	36	440	0.0109
7	0	30	100	0	400	37	500	0.0267
0	20	20	0	300	250	40	550	0.0672
7	9	30	100	90	400	46	590	0.0073
7	20	20	100	300	250	47	650	0.0394
0	20	30	0	300	400	50	700	0.1162
7	20	30	100	300	400	57	800	0.1938
7	20	40	100	300	500	67	900	0.0834

<i>Panel B</i>				
Year	Goal#	Cost	Utility	Probability
5	1	0	0	0.3579
	1	7	100	0.6421
	2	0	0	0.1019
	2	9	90	0.0194
	2	20	300	0.8787
	3	0	0	0.4551
	3	20	250	0.1175
	3	30	400	0.3440
	3	40	500	0.0834
10	1	0	0	0.4187
	1	50	500	0.2137
	1	90	1000	0.3676

that is mildly close to that of the tuition goal.

From Panel B of Table 6, we also see that the probability of attaining the tuition goal decreases each time we jump forward a year, due to the tuition cost going up while the utility remains constant. However, as we change semesters within years 3, 4, and 5, the probability of attaining the tuition goal increases, since the later semester has an extra half a year for the investor's portfolio to grow. In contrast, as we change semesters within year 6, the probability decreases, instead of increases. This makes sense as well:

Since the tuition goals at year 6 and at year 6.5 have the same utility, the goal at year 6 should always be taken if the investor has the money, because that guarantees a utility of at least 1000, whereas forgoing this goal guarantees a future utility of at most 1000 from the tuition goal at year 6.5. Since it is possible that fulfilling the tuition goal at year 6 means not having enough money to pay the last semester of tuition at year 6.5, we expect the probability of fulfilling the tuition goal at year 6.5 to be lower than the probability of fulfilling the tuition goal at year 6, just as we observe.

Table 6: Car and tuition goals example from Subsection 2.2. All cost and wealth numbers are in thousands of dollars. The initial wealth is $W(0) = \$100$. The format for this table is the same format used for Table 5. Goal 1 is for tuition, except at year 4, when Goal 1 is for the car and Goal 2 is for tuition.

<i>Panel A: Year 4 (t=8)</i>						
Cost		Utility		Total		
Goal 1	Goal 2	Goal 1	Goal 2	Cost	Utility	Probability
0	0	0	0	0	0	0.2309
0	20.407	0	1000	20.407	1000	0.7633
28	20.407	80	1000	48.407	1080	0.0014
32	20.407	125	1000	52.407	1125	0.0021
50	20.407	300	1000	70.407	1300	0.0023

<i>Panel B</i>				
Year	Goal#	Cost	Utility	Probability
3	1	0.000	0	0.1412
	1	18.895	1000	0.8588
3.5	1	0.000	0	0.0076
	1	18.895	1000	0.9924
4	1	0.000	0	0.9942
	1	28.000	80	0.0014
	1	32.000	125	0.0021
	1	50.000	300	0.0023
	2	0.000	0	0.2309
	2	20.407	1000	0.7691
4.5	1	0.000	0	0.0765
	1	20.407	1000	0.9235
5	1	0.000	0	0.2235
	1	22.039	1000	0.7765
5.5	1	0.000	0	0.1314
	1	22.039	1000	0.8686
6	1	0.000	0	0.2837
	1	23.802	1000	0.7163
6.5	1	0.000	0	0.5500
	1	23.802	1000	0.4500

4.4 Optimal Long Range Financial Planning for Retirement

This section shows that target-date funds materially underperform the GBWM algorithm.

We consider an individual investor who is 35 years old now (at $t = 0$). The investor is planning for their retirement at age 70. We will assume the investor currently has \$100 in their retirement account (meaning one hundred thousand dollars, since monetary figures in this example are in thousands of (after-tax) dollars). Assuming a rate of inflation of 3%, the investor intends to contribute \$10 real $t = 0$ dollars to their retirement plan, that is, $\$10 \cdot 1.03^t$, every year until they stop working at age 69 ($t = 34$). Starting at age 70, they are projecting they will annually collect $\$80 \cdot 1.03^{(t - 35)}$ dollars, that is \$80 real $t = 35$ dollars from Social Security.

This investor has three (all-or-nothing) goals:

- Goal 1: Purchase a \$2500 (that is 2.5 million dollar) annuity at the age of 70, which should initially (at age 70) annually generate about \$60 (thousand) in $t = 0$ dollars,³ but then does not grow with inflation in later years.
- Goal 2: Purchase a \$3000 annuity at age 85 to cover potential new medical and long-term care costs. This should initially (at age 85) annually generate about \$90 (thousand) in $t = 0$ dollars, but, as before, does not grow with inflation after that.
- Goal 3: Give \$4000 at age 95 to grandchildren to help with the cost of a college education or with a downpayment on a new house. This amount corresponds to \$680 (thousand) in $t = 0$ dollars.

We compare two methods of investing for these future goals. The first method determines the investment portfolio by using a typical target date fund approach. The approach uses the same three index funds (U.S. Bond, International Stocks, and U.S. Stocks) that we have used in the other examples. The proportions for the three funds are given by the glide path shown in Table 7.⁴ As discussed earlier, this approach, which has become progressively popular with advisors and clients, cannot prioritize goals, so it will take each goal if there are sufficient funds. Further, the glide path approach chooses an asset mix strictly as a function of time, and therefore the assets chosen have no relation to how well the investor’s portfolio is performing or the goals the investor looks to attain. It is a one-size-fits-all approach.

The second method uses the dynamic programming approach. We will assign a utility of $u = 4$ to Goal 1, $u = 2$ to Goal 2, and $u = 1$ to Goal 3. These utility values were

³This Goal 1 annuity approximation and the annuity approximation in Goal 2 were generated using https://www.schwab.com/public/schwab/investing/accounts_products/investment/annuities/income_annuity/fixed_income_annuity_calculator. Optimization with annuities has been explored in a series of excellent papers: Konicz and Mulvey (2013); Konicz and Mulvey (2015); Konicz et al. (2016).

⁴The term “glide path” is used in the standard sense for target-date funds, see: <https://www.investopedia.com/terms/g/glide-path.asp>. The definition is stated as follows: “Glide path refers to a formula that defines the asset allocation mix of a target-date fund, based on the number of years to the target date. The glide path creates an asset allocation that typically becomes more conservative (i.e., includes more fixed-income assets and fewer equities) as a fund gets closer to the target date.”

Table 7: The target date fund glide path used in Subsection 4.4

Age range	35–39	40–44	45–49	50–54	55–59	60–64	65–69	70–74	75–95
1. U.S. Stock	63%	63%	60%	55%	50%	45%	38%	29%	21%
2. International Stock	27%	27%	25%	23%	20%	18%	16%	12%	9%
3. U.S. Bond	10%	10%	15%	22%	30%	37%	46%	59%	70%

chosen because they correspond to an optimal strategy that fulfills each goal if there are sufficient funds. In other words, the optimal goal taking strategy is chosen to match what the target date fund’s goal taking strategy must be. So with the goal taking strategy being the same, the dynamic programming approach is only able to take advantage of the optimal investment portfolio strategy it provides, as opposed to the optimal goals taking strategy.

How much of a difference does this optimal investment portfolio strategy make? Table 8 shows that the difference is considerable. At age 70, the dynamic programming approach will have the \$2500 needed to attain the first goal 78.2% of the time, while the target date fund approach will have sufficient funds only 63.9% of the time. This difference widens when the investor reaches the two remaining goals. By the end of the third goal at age 95, the investor who opts for the target date fund approach has a 55.1% chance of going bankrupt, while the investor who opts for the dynamic programming approach reduces that chance by 26.4 percentage points to 28.7%.

Table 8: Chance of remaining solvent after paying each goal

	Goal 1	Goal 2	Goal 3
Target Date Fund approach	63.9%	55.9%	44.9%
Dynamic Programming approach	78.2%	74.6%	71.3%

This example shows that the beneficial effects of using the dynamic programming approach are far from trivial for investors. In fact, to attain the 71.3% chance of remaining solvent at age 95 that dynamic programming has, the investor using the target date fund approach would have needed to start at age 35 with \$280 (thousand) instead of \$100 (thousand).

4.5 Optimal Long Range Financial Planning with Numerous Lifetime Goals

Next we consider a couple, both 35 years old, with one 5 year old child. The couple is looking to create a long term goals-based investment plan over the next 60 years, until they are 95. (So $T = 61$.) Their intent is to retire and start taking Social Security at age 70. As in the previous examples, all of the goals and costs here are in thousands of (after-tax) dollars.

The income sources (infusions) for the couple are:

1. Salary income: A combined annual salary of \$75, projected to increase annually

at 4% until retirement at age 69 ($t = 34$). This includes money being saved for retirement.

2. Retirement income: Social Security is assumed to pay \$120 at age 70⁵ ($t = 35$), which corresponds to around \$43 present day dollars, assuming a rate of inflation of 3%. We will also assume the rate of inflation continues at 3% for Social Security payouts between $t = 35$ and $t = 60$.

The couple's goals are:

1. Tier 1 (highest priority) goals

- (a) Mortgage: Assume the couple needs to pay a fixed annual rate of \$10 for the next 25 years ($t = 1$ through $t = 25$) to pay off what remains from a 30 year fixed mortgage.
- (b) Property tax: Assume the annual cost is \$6, which goes up every year by 2%.
- (c) Long term care insurance: Starting at age 50 ($t = 15$), at a cost of \$7, which then goes up by 4% every year.
- (d) Medical expenses: Assume these are insignificant until the age of 75, after which they are approximated to be \$8 initially with an inflation rate of 10%. The high inflation rate reflects both progressively higher costs of medical care, as well as progressively higher needs for care as the couple ages.
- (e) Everyday expenses: Assume these start at \$60 per year and go up at a 3% rate of inflation. Further, if necessary, these can be trimmed to \$50 per year by cutting costs that are not as crucial.

2. Tier 2 goals

- (a) Orthodontia for the couple's child: Assume this costs \$3 each year when their child is 11, 12, and 13; that is, $t = 6, 7$, and 8.
- (b) College tuition for the couple's child: Assume the annual cost is \$40 at $t = 13$, when the child is 18, and then goes up by 8% in each of the next three years. A partial goal that starts at \$25 and then goes up by 8% is also available.
- (c) Cars: Assume the couple would like to purchase a new car every five years at a cost of \$32. They may also opt for the partial goal of a used car for \$22.

3. Tier 3 goals

- (a) Remodeling the house at $t = 16$: A nice remodel will cost \$50. Less nice remodels at costs of \$40 and \$25 are also available.
- (b) Wedding expenses for the couple's child: Assume this costs \$70, with the couple approximating that their child will marry at age 28; that is, $t = 23$. A less expensive wedding costs \$55.

⁵Approximated using Social Security's quick benefit estimate calculator at <https://www.ssa.gov/cgi-bin/benefit6.cgi>.

4. Tier 4 (lowest priority) goals

- (a) Private high school for the couple’s child: Assume the annual cost is \$25 at $t = 9$, when the child is 14, and then goes up by 5% in each of the next three years.
- (b) A fancy trip: Assume the couple wishes to take a nice trip every 10 years at a cost of \$15 at $t = 10$, which increases at the annual inflation rate of 3% thereafter.
- (c) Philanthropic gifts: Initially, \$5 every year or a partial goal of \$2.5. This cost goes up at a 3% annual rate of inflation.

That is, in total over the 60 year timeframe, the couple is considering 301 full goals and 138 partial goals.

We assume the initial wealth, $W(0)$, is \$100, therefore $W_{\max} = \$18,000$ (that is, \$18 million). We have selected $i_{\max} = 1221$ for the number of nodes in the wealth grid between $W_{\min} = 1$ and W_{\max} . Because of the significant increase in the number of goals, the horizon time T , and the number of wealth grid points, the runtime increases from under 2 seconds needed in previous examples to 17 seconds. The utility values were adjusted so that the optimized probabilities correspond to the priority of the tier assigned to each goal by the couple. It should be noted that, unsurprisingly, this requires assigning larger utility values to the more costly goals within a given tier. On occasion, assigning lower utility values to later goals within a tier was also useful. Table 9 gives the final utility values assigned to the goals and shows the corresponding optimal probabilities of reaching the partial and full goals.

The optimal expected total utility is 574,239 out of a maximum possible of 577,055, so the expected utility fraction is $574,239/577,055 = 0.9951$. It is reasonable to think at first that such a high fraction means a high likelihood of attaining all the couple’s goals, but a quick overview of Table 9 shows that this is not the case. The high fraction simply means the probability of attaining the goals with high utility values is quite high. For example, if the utility assigned to a single goal were made extremely high and the probability of attaining that goal were high, the expected utility fraction would be close to one, even if the chances were quite low of attaining any of the other 300 goals.

In order to assess how much of a difference *fat-tailed* distributions make to the final results, we replaced the normal distribution with a t-distribution with 5 degrees of freedom. The total expected utility becomes 573,844, which corresponds to an expected utility fraction of 0.9944, as compared to the total expected utility of 574,239, which corresponds to the fraction 0.9951 just shown above for the normal distribution. So the impact on the outcomes is small, however, it does mean that some goals will not be met in comparison to the normal distribution case. We further reduced the degrees of freedom to 2, and obtained an expected utility of 572,944 (0.9929). The fat-tails make a difference and in order to assess the effect of these small differences in utility, see the next Section 4.5.1, which compares a good path versus a poor one, where differences in expected utility can be seen as differences in goals met. We note, as mentioned in Subsection 2.1, that working with any desired Markovian evolution model with deterministic parameters is straight-forward, so replacing our geometric Brownian motion model with any theoretical

Table 9: Optimal probability ranges for the couple’s goals

Goal	Goal Years (t)	Initial Cost	Inflation Rate	Utility	Range for optimal probability of fulfillment
1a: Mortgage	1–25	\$10	0%	500	> 0.99
1b: Prop. tax	1–60	\$6	2%	1000	> 0.99
1c: LTC Insur.	15–60	\$7	4%	1035	> 0.99
1d: Med exp.	40–60	\$8	10%	1060	> 0.99
1e: Everyday	1–60		3%		> 0.99
(full)	1–60	\$60		7000	(> 0.97)
(partial)	1–60	\$50		5950	(< 0.03)
2a: Orthodon.	6–8	\$3	0%	100	> 0.99
2b: College	13–16		8%		> 0.99
(full)	13–16	\$40		900	(0.56–0.60)
(partial)	13–16	\$25		750	(0.40–0.43)
2c: Cars	5,10,...,55		0%		> 0.97
(full)	5,10,...,55	\$32		600–710	(0.57–0.98)
(partial)	5,10,...,55	\$22		400–550	(0–0.41)
3a: Remodel	16				0.58
(full)	16	\$50		300	(0.23)
(partial 1)	16	\$40		250	(0)
(partial 2)	16	\$25		200	(0.35)
3b: Wedding	23				0.67
(full)	23	\$70		400	(0.38)
(partial)	23	\$55		350	(0.29)
4a: High Sch.	9–12	\$25	5%	100	0.04–0.07
4b: Trips	10, 20,...,50	\$15	3%	50–100	0.20–0.78
4c: Philanth.	1–60		3%		0–0.97
(full)	1–60	\$5		40	(0–0.92)
(partial)	1–60	\$2.5		30	(0–0.45)

fat-tailed distribution, for example, requires minimal changes to the algorithm. Another approach is to use fat-tails from the historical record using the simulation approaches suggested in [Geyer et al. \(2014\)](#) and [Consiglio et al. \(2016\)](#).

Finally, we note that any evolution model may have problems being accurate with long-term predictions. We therefore expect that investors will periodically rerun the algorithm presented here with updates to the evolutionary model, as well as any changes in their goal preferences. It is also possible to make the parameters in the Markovian evolution process stochastic to address potential changes over time, but this comes at the potentially significant cost of adding an additional state space variable for each parameter, which can slow the algorithm considerably, and creates the new problem of trying to accurately model the stochastic process governing each parameter’s evolution.

Table 9 clearly shows how the algorithm uses the utility weights to prioritize the tiered goals. Each Tier 1 goal has a greater than 99% chance of being fulfilled and a greater than 97% chance of being fully fulfilled. Each Tier 2 goal has a greater than 97% chance of being fulfilled and a greater than 50% chance of being fully fulfilled. Each Tier 3 goal has a greater than 50% chance of begin fulfilled, and the probability of attaining the Tier 4 goals varies from 0 to 98%, largely depending on the investors’ age. That

is, our optimized dynamic programming results for this couple show that by following the algorithm’s advice for making optimal, prudent decisions in both goals taking and investing, they look to have a rosy financial future.

4.5.1 Examining a few sample paths with multiple goals

Whereas the discussion above encompasses the entire distribution of outcomes, it is instructive to focus on a couple of sample paths to see how the algorithm decides which goals to take and which to forgo. This path dependence is illustrated by considering two paths, one where the portfolio has good returns, and another one where the portfolio’s returns are comparatively poor. These are generated by simulating the wealth evolution of the portfolio while applying the optimal controls for portfolio allocation and goal choices along each path. The results are shown in Figure 4.

In the top panel of Figure 4, we see that the total utility achieved along the good path is 576,485 out of a possible total of 577,055, whereas total utility along the poor path is 574,540. From the middle panel of Figure 4, we note the following differences between the two cases: (1) The portfolio in the poor case underperforms the good case in the early years, leading to either partially or completely reduced philanthropy (a Tier 4 goal) in years 7–37. (2) Another Tier 4 goal (High School) is not met in year 9 in both good and poor cases, but had to be forgone in the poor case also in years 10-12, but not in the good case. (3) Trips (Tier 4) are also passed up in year 10 and in year 20 in the poor case. (4) In the poor case, the college goal could only be partially met in years 14-16 when the portfolio performed poorly. (5) The remodel goal in year 16 is also passed up completely in the poor case. (6) In years 20, 25, 30 the car goal can only be partially taken in the poor case. (7) In year 23, the wedding could not be funded in the poor case. (8) All this early belt-tightening pays off for the investor with the poor case, since it allows all goals to be met after year 37.


Overall, from the bottom panel of Figure 4, we see that in the early years, the poor portfolio does not perform that well, leading to passing on goals, especially in the second decade. After that, the poor portfolio catches up and all later goals are taken. It is precisely because the algorithm is taking the future into account, that it passes on some goals in the early years in order to enable taking later goals.

4.5.2 Comparison to Monte Carlo Methods

Recall from the introduction that Monte Carlo methods require a fixed strategy, like staying with one of the 15 portfolios that are available to the dynamic programming strategy or, as we saw in Subsection 4.4, applying a target date fund strategy. Further, Monte Carlo methods only look to fully attain each goal, as decisions like fulfilling or not fulfilling goals or partially fulfilling goals will quickly lead to a crippling exponential blow up in Monte Carlo computational time.

Separately applying each of the 15 portfolios available to the Monte Carlo method to the example above, we find that the most aggressive, portfolio 15, has the highest probability of fulfilling all of the couple’s goals. However, this optimal probability is only

Top panel: Two possible paths for portfolio wealth under good and poor scenarios.



Middle panel: Total utility achieved for each goal under good and poor scenarios.

Goal	Good Scenario	Poor Scenario
Mortgage payments	12500	12500
Property taxes	60000	60000
Long-term care insurance	47610	47610
Medical expenses	22260	22260
Everyday expenses	420000	420000
Orthodontia	300	300
College	3600	3150
Cars	6960	6760
Remodel	300	0
Wedding	400	0
High school	300	0
Trips	275	100
Philanthropy	1980	920
Total	576485	574540

Bottom panel: Additional goals taken in the good scenario versus the poor one.




Figure 4: We examine two scenarios, one for a path with good returns, the other for a path with poor returns. These are shown in the top panel. The table in the middle panel shows the total utility achieved for each goal, where we see that Tier 1 and Tier 2 goals are achieved, because the algorithm optimally forgoes Tier 3 and Tier 4 goals as needed. The bottom panel displays the additional goals taken on the good path versus the poor one.

7%, and, indeed, its probability of even fulfilling all the goals through year $t = 16$, when the couple's child leaves college, is only 11%. The standard financial advice in a situation like this is to remove lower tier goals. If we remove the Tier 4 goals, portfolio 8 becomes optimal, which corresponds to a 48% chance of attaining the remaining goals. If we remove the Tier 3 and 4 goals, portfolio 4 becomes optimal, which corresponds to a 68% chance of attaining the remaining goals. If we drop all Tier 2, 3, and 4 goals, portfolio 1 becomes optimal, which is essentially all bonds, and yields a greater than 99% chance of attaining the remaining goals.

The practical impact on the couple if they are presented these Monte Carlo based results is clear: Believing that their financial future is quite perilous, they would give up on attaining their Tier 3 and 4 goals and either live with the belief that there was an over 30% chance they would eventually go broke attempting to attain their Tier 1 and Tier 2 goals, or they might jettison some or all of the full goals in Tier 1 and Tier 2 in favor of what were previously partial goals. In contrast, by using our dynamic programming approach, we see that the probability of the couple being able to fulfill their Tier 1 and Tier 2 goals is actually excellent, and there is also a good chance of fulfilling some of their Tier 3 and Tier 4 goals as well.

5 Concluding Discussion

There is a great deal of interest in how to best apply a goals-based wealth management approach to investors hoping to attain as many of their goals as possible, weighted by their importance. We have shown that dynamic programming methods can answer three key questions for investors pursuing multiple goals that Monte Carlo methods cannot answer:

1. When an investor has limited means, they must choose whether or not to fulfill or forgo each of their goals as they progress through time. Given the importance of each of the goals to the investor, we have shown how to optimally determine whether or not to fulfill or forgo each goal. Implicit in this analysis is the idea that portfolios may be chosen with any risk-return tradeoff the investor prefers, so is agnostic to risk measures (see [Szegö \(2002\)](#)), but also introduces is cognizant of behavioral risk, i.e., the probability of not achieving a goal.
2. We have shown how to optimally, dynamically determine which investment portfolio the investor should use at any time based on their portfolio worth at that time.
3. Under the optimized goals taking and investment portfolio selection strategy in the preceding two points, we have shown how to determine the corresponding probabilities for attaining each of the investor's goals. Should the investor feel that these probabilities do not reflect their preferences, they may change the weights of the goals and rerun the algorithm until the generated probabilities match the investor's priorities.

We have also shown how the multiple goals method can easily accommodate concurrent goals (that is, multiple goals that occur at the same time) and partial goals (when

the investor is open to options that cost less than the full goal, understanding that these partial goals will make the investor less happy than the full option). On a simple desktop or laptop computer, we are able to compute the optimal solution for a small number of multiple goals in under 2 seconds. For a much more involved multiple goals scenario over a decades-long horizon, the computation still runs in 17 seconds.

Our focus in this paper has been to explain the large picture of how an investor can achieve a small or large number of purchasing goals with a significantly greater probability than methods typically used in the wealth management industry. Of course, there are limitations and a number of potential additional features that we have not pursued in this paper, which we leave to future work. These include:

1. Allowing mandatory goals, which must to be taken even if it means bankrupting the portfolio. This however, may not apply to most investors.
2. Allowing for stochastic (unexpected) expenses along the path. This is best handled using stochastic programming approaches mentioned in the introduction.
3. While we have not handled infusions or withdrawals that are correlated with wealth levels, we can easily extend the model to these. For example, at low wealth levels we may infuse money by belt-tightening elsewhere, whereas at high levels we may choose to make optional gifts.
4. We may extend the algorithm to also optimize taxes using forward simulations. Since the tax code is complicated, this can only be handled through forward simulations, suggesting backward recursion for portfolio optimization combined with a forward traversal of the portfolio that accounts for taxes. These backward and forward steps may be iterated a few times until a stable solution is achieved. We leave this for a large-scale follow-up paper.
5. The model does not account for investor mortality but can be easily extended to handle this in the backward recursion. One may imagine that the objective function can be adapted to maximize the number of goals reached while remaining solvent through the lifetime of the investor. If we wish to handle mortality risk for a couple, then we need to account for death of each of the investors, which doubles the size of the state space.
6. As discussed in the introduction, our approach does not allow for optimally deciding to postpone a goal. The approach in the paper requires a take it or leave it approach. Stochastic programming or reinforcement learning methods are better equipped to handle such phenomena.

The algorithm presented in this paper is adept at handling a very large action space, which allows it to address large numbers of goals and investment portfolio options quickly and efficiently, but not a very large state space, since, like all dynamic programming methods, the state space is subject to the so-called “curse of dimensionality.” Alternative approaches include the backward-forward method mentioned earlier, stochastic programming, and reinforcement learning, however, these methods impose constraints on the action space. Understanding the trade-off between the restrictions on the action

space versus the restrictions on the dimensionality of the state space is key to determining which method is best to apply in specific applications.

In summary, we have prescribed a goals-based wealth management plan for multiple goals that may be implemented in a computationally efficient manner over a very long horizon. The goals may occur at different times or concurrently. The goals may be fulfilled partially, if the investor wishes this possibility, or fully. The algorithm maximizes an investor's expected utility, via the optimal exercise of goals and the optimal selection of investment portfolios. These optimal decisions are determined using dynamic programming, which optimally balances the trade-off between cost and utility for each goal in the plan. The algorithm provides the investor with the probabilities of achieving each goal so that the investor may adjust the plan in an iterative manner. Changes in these probabilities may be tracked over time as an additional performance metric to complement traditional metrics that compare risk-adjusted returns to a benchmark. Thus, this work offers a comprehensive approach to achieving multiple goals in wealth management, thereby improving on (i) approaches that handle multiple goals in separate mental account portfolios and (ii) approaches that use Monte Carlo simulation.

References

- Alexander, G. J., A. M. Baptista, and S. Yan (2017, July). Portfolio Selection with Mental Accounts and Estimation Risk. SSRN Scholarly Paper ID 3001158, Social Science Research Network, Rochester, NY.
- Bellman, R. (1952, August). On the Theory of Dynamic Programming. Proceedings of the National Academy of Sciences 38(8), 716–719.
- Bergerson, A., L. Kotlikoff, T. Matina, and A. Seychuk (2016). Rethinking Portfolio Analysis: Living Standard Risk and Reward Over the Life Cycle.
- Browne, S. (1997, May). Survival and Growth with a Liability: Optimal Portfolio Strategies in Continuous Time. Mathematics of Operations Research 22(2), 468–493.
- Browne, S. (1999). Reaching Goals by a Deadline: Digital Options and Continuous-Time Active Portfolio Management. Advances in Applied Probability 31(2), 551–577.
- Brunel, J. (2015). Goals-Based Wealth Management: An Integrated and Practical Approach to Changing the Structure of Wealth Advisory Practices. New York: Wiley.
- Chhabra, A. B. (2005, January). Beyond Markowitz: A Comprehensive Wealth Allocation Framework for Individual Investors. The Journal of Wealth Management 7(4), 8–34.
- Consigli, G. and M. A. H. Dempster (1998, June). Dynamic stochastic programming for asset-liability management. Annals of Operations Research 81(0), 131–162.
- Consiglio, A., A. Carollo, and S. A. Zenios (2016). A parsimonious model for generating arbitrage-free scenario trees. Quantitative Finance 16(2), 201–212. Publisher: Taylor & Francis Journals.

- Consiglio, A., F. Cocco, and S. Zenios (2007). Scenario optimization asset and liability modelling for individual investors. Annals of Operations Research 152(1), 167–191. Publisher: Springer.
- Consiglio, A., F. Cocco, and S. A. Zenios (2004). www.Personal_asset_allocation.Interfaces 34(4), 287–302. Publisher: INFORMS.
- Dantzig, G. B. and G. Infanger (1993, December). Multi-stage stochastic linear programs for portfolio optimization. Annals of Operations Research 45(1), 59–76.
- Das, S. R., H. Markowitz, H., J. Scheid, and M. Statman (2010). Portfolio Optimization with Mental Accounts. Journal of Financial and Quantitative Analysis 45(2), 311–334.
- Das, S. R., D. Ostrov, A. Radhakrishnan, and D. Srivastav (2018). Goals-Based Wealth Management: A New Approach. Journal of Investment Management 16(3), 1–27.
- Das, S. R., D. Ostrov, A. Radhakrishnan, and D. Srivastav (2019). A Dynamic Approach to Goals-Based Wealth Management. Computational Management Science *forthcoming*, 1–28.
- Deguest, R., L. Martellini, V. Milhau, A. Suri, and H. Wang (2015). Introducing a Comprehensive Allocation Framework for Goals-Based Wealth Management. Working paper, EDHEC Business School.
- Dempster, M. and E. Medova (2011). Planning for Retirement: Asset Liability Management for Individuals. In G. Mitra and K. Schwaiger (Eds.), Asset and Liability Management Handbook, pp. 409–432. London: Palgrave Macmillan UK.
- Dempster, M. a. H., M. Germano, E. A. Medova, and M. Villaverde (2003, April). Global Asset Liability Management. British Actuarial Journal 9(1), 137–195. Publisher: Cambridge University Press.
- Geyer, A., M. Hanke, and A. Weissensteiner (2014, August). No-Arbitrage ROM simulation. Journal of Economic Dynamics and Control 45, 66–79.
- Infanger, G. (2008, January). Chapter 5 - Dynamic asset allocation strategies using a stochastic dynamic programming approach. In S. A. Zenios and W. T. Ziemba (Eds.), Handbook of Asset and Liability Management, pp. 199–251. San Diego: North-Holland.
- Kahneman, D. and A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. Econometrica 47(2), 263–291.
- Kim, W. C., D.-G. Kwon, Y. Lee, J. H. Kim, and C. Lin (2020, March). Personalized goal-based investing via multi-stage stochastic goal programming. Quantitative Finance 20(3), 515–526. Publisher: Routledge eprint: <https://doi.org/10.1080/14697688.2019.1662079>.
- Konicz, A. K. and J. M. Mulvey (2013, October). Applying a Stochastic Financial Planning System to an Individual: Immediate or Deferred Life Annuities? The Journal of Retirement 1(2), 46–60. Publisher: Institutional Investor Journals Umbrella Section: Article.

- Konicz, A. K. and J. M. Mulvey (2015, May). Optimal savings management for individuals with defined contribution pension plans. European Journal of Operational Research 243(1), 233–247.
- Konicz, A. K., D. Pisinger, K. M. Rasmussen, and M. Steffensen (2015, July). A combined stochastic programming and optimal control approach to personal finance and pensions. OR Spectrum 37(3), 583–616.
- Konicz, A. K., D. Pisinger, and A. Weissensteiner (2015, July). Optimal annuity portfolio under inflation risk. Computational Management Science 12(3), 461–488.
- Konicz, A. K., D. Pisinger, and A. Weissensteiner (2016, February). Optimal retirement planning with a focus on single and joint life annuities. Quantitative Finance 16(2), 275–295. Publisher: Routledge eprint: <https://doi.org/10.1080/14697688.2015.1114361>.
- Markowitz, H. H. (1952). Portfolio Selection. Journal of Finance 6, 77–91.
- Merton, R. (1969). Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. The Review of Economics and Statistics 51(3), 247–57.
- Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. Journal of Economic Theory 3(4), 373–413.
- Mulvey, J. M. and B. Shetty (2004, January). Financial planning via multi-stage stochastic optimization. Computers & Operations Research 31(1), 1–20.
- Mulvey, J. M. and H. Vladimirou (1992). Stochastic Network Programming for Financial Planning Problems. Management Science 38(11), 1642–1664. Publisher: INFORMS.
- Muralidhar, A. (2018, December). The F-Utility of Wealth: It’s All Relative. SSRN Scholarly Paper ID 3329917, Social Science Research Network, Rochester, NY.
- Nevins, D. (2004, January). Goals-Based Investing: Integrating Traditional and Behavioral Finance. The Journal of Wealth Management 6(4), 8–23.
- Pakizer, S. (2017, July). Goals Based Investing for the Middle Class: Can Behavioral and Traditional Financial Policies Simplify Investing and Maximize Wealth? SSRN Scholarly Paper ID 3036231, Social Science Research Network, Rochester, NY.
- Parker, F. J. (2020, April). Allocation of Wealth Both Within and Across Goals: *A Practitioner’s Guide*. The Journal of Wealth Management 23(1), 8–21.
- Parker, F. J. (2021, January). A Goals-Based Theory of Utility. Journal of Behavioral Finance 22(1), 10–25. Publisher: Routledge eprint: <https://doi.org/10.1080/15427560.2020.1716359>.
- Shefrin, H. (2008, June). A Behavioral Approach to Asset Pricing (2 edition ed.). Amsterdam ; Boston: Academic Press.
- Shefrin, H. and M. Statman (2000). Behavioral Portfolio Theory. The Journal of Financial and Quantitative Analysis 35(2), 127–151.

- Simsek, K. D., M. J. Kim, W. C. Kim, and J. M. Mulvey (2018, January). Optimal Longevity Risk Management in the Retirement Stage of the Life Cycle. The Journal of Retirement 5(3), 73–92. Publisher: Institutional Investor Journals Umbrella Section: Article.
- Szegö, G. (2002, July). Measures of risk. Journal of Banking & Finance 26(7), 1253–1272.
- Szegö, G. P. (2014, May). Portfolio Theory: With Application to Bank Asset Management. Academic Press. Google-Books-ID: KqjBQAAQBAJ.
- Thaler, R. (1985, August). Mental Accounting and Consumer Choice. Marketing Science 4(3), 199–214.
- Topaloglou, N., H. Vladimirov, and S. A. Zenios (2008, March). A dynamic stochastic programming model for international portfolio management. European Journal of Operational Research 185(3), 1501–1524.
- Wang, H., A. Suri, D. Laster, and H. Almadi (2011, April). Portfolio Selection in Goals-Based Wealth Management. The Journal of Wealth Management 14(1), 55–65.

6 Appendix: The Effect of Changing Conditions on the Multiple All-or-Nothing Example from Section 4.2

6.1 The Effect of Changing the Utility Assigned to a Goal

The investor may look at the probabilities generated in Table 3 and decide that some are lower than they want. For example, let’s say the investor feels that a 3.08% chance of attaining Goal 4 is too low. They may therefore decide to increase the utility assigned to Goal 4. Results generated from increasing the utility for Goal 4 are shown in Table 10. The table shows that increasing the utility assigned to Goal 4 increases the probability that Goal 4 is attained, but, in general, it also decreases the probability that the other goals are attained.

Note that this is not always the case, however. For example, in Table 10, we see that the probability of attaining Goal 3 *increases* when the utility for Goal 4 increases from 1500 to 2000. This is because the utility increase encourages more aggressive portfolios early on, which increases the chance that attaining Goal 3 at $t = 10$ is both possible and worthwhile, especially if Goal 4, which is at $t = 11$, remains unattainable, since the \$80 (thousand) Goal 4 cost is much higher than the \$15 (thousand) Goal 3 cost.

Because Goal 2 generates a high utility, 2500, at the low cost of \$17 (thousand), it remains heavily desirable until the utility assigned to Goal 4 increases past 5000. At that point, it becomes progressively more important to consider abandoning Goal 2, so as to maximize the chance of attaining Goal 4.

Table 10: The effect of the investor increasing the utility value assigned to Goal 4 upon the base case specified in Table 3, reprised at the top of the table. This increases the chances of attaining Goal 4, but generally decreases the chance of attaining the other goals.

	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
Goal time, t (in years)	5	8	10	11	17	22	24
Goal cost, c (in 1000s of \$)	25	17	15	80	50	60	130
Goal utility, u	1000	2500	500	1500	300	3000	2000

Utility, u , for Goal 4	Probability of fulfilling goal						
	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
1500	0.0446	0.9871	0.2077	0.0316	0.0192	0.7569	0.2373
2000	0.0309	0.9860	0.2244	0.0596	0.0158	0.7526	0.2261
3000	0.0073	0.9854	0.1327	0.1314	0.0106	0.7329	0.2022
5000	0.0021	0.9853	0.0342	0.2341	0.0049	0.6561	0.1438
10000	0.0002	0.6972	0.0098	0.3857	0.0021	0.5234	0.0839
100000	0.0000	0.1258	0.0002	0.4776	0.0025	0.4997	0.1082

Note also that when the utility for Goal 4 increases from 10,000 to 100,000, the algorithm adopts maximally aggressive portfolios early on. This increased volatility increases the chance of extremely high or low wealth values later on when the investor is looking to attain Goals 6 and 7. Attaining Goal 6 has priority over Goal 7 because it is much less expensive and has a significantly higher utility. Therefore, we see the approximate likelihood of attaining neither of these final goals due to low wealth increasing from $1 - 0.5234 = 0.4766$ to $1 - 0.4997 = 0.5003$ and the approximate likelihood of attaining both goals due to high wealth increasing from 0.0839 to 0.1082. This last increase corresponds to the numbers in Table 10 for the probability of attaining Goal 7, since Goal 6, with its higher priority, is generally attained whenever Goal 7 is attained. This increase in the probability of attaining Goal 7 is another example of where increasing the utility assigned to Goal 4 can increase, instead of decrease, the probability of attaining a goal other than Goal 4.

The investor, as much as they desire, can proceed to increase (or decrease) the assigned utilities for each goal and see the effect this has on the optimal probabilities of attaining their goals. At the end of the day, they will have an accurate understanding of the trade-offs between all of their goals and the limitations of their current portfolio under the set of available investment portfolios, when optimally run.

6.2 The Effect of Changing the Initial Wealth, $W(0)$

Should the investor wish to increase their chances of obtaining all of their goals, they may decide to increase their initial investment, $W(0)$. Figure 5 shows how the optimally achievable total expected utility fraction depends upon $W(0)$. The red dot in the figure represents the base case from Table 3, where $W(0) = \$30$ and the expected utility fraction is 0.5018. Should the investor increase their original investment so that $W(0) = \$50$

instead of \$30, then, from the figure, we can see that the expected utility fraction will increase from approximately $1/2$ to approximately $2/3$. As $W(0)$ gets larger, the expected utility fraction approaches 1, of course, and, similarly, as $W(0)$ gets small, the expected utility fraction goes to 0. We note that because W_{\max} in the algorithm is determined from the value of $W(0)$, to generate Figure 5 we had to make certain that W_{\max} was sufficiently large to accommodate all the values of $W(0)$ used in the figure.




Figure 5: The effect of varying the initial wealth, $W(0)$, on attaining the investor’s goals. Information for the seven goals used here is presented in Table 3. The initial wealth base case of \$30 corresponds to an optimally achievable total expected utility fraction of 0.5018, as represented by the red dot.

6.3 The Effect of Infusions

Should the investor not have the funds or the desire to increase the initial investment, they may decide instead to commit to future cash infusions, $I(t)$. The infusions can be for any desired amount. Typical infusion examples include being constant ($I(t) = k$), adjusted for inflation ($I(t) = k(1 + r)^t$), or a one-time infusion ($I(t_0) = k$ while $I(t) = 0$ if $t \notin t_0$).

Panel A of Table 11 shows the beneficial effect on the base case in Table 3 of having constant annual infusions, starting at the beginning of year 1 and ending at the beginning of year 24. The first row in the panel has no infusions and is therefore the same as the base case in Table 3. As we increase the constant infusion amount, we see an increase not only in the expected utility fraction, as must be the case, but also in the probabilities of attaining each individual goal. We note that even small infusions have a substantial impact on achieving goals, making a strong mathematical argument in support of saving and investment. Further, the algorithm allows an investor to understand in clear, concrete terms how setting aside money on a regular basis directly impacts their (now quantifiable) chances of attaining their goals.

Panel B of Table 11 shows the effect on the base case in Table 3 of having one-time infusions. We note that the first and fourth rows correspond to infusions at $t = 0$, which is the equivalent of boosting $W(0)$ from its base case value of \$30. Therefore, from these rows we see that $W(0) = \$30 + \$30 = \$60$ induces an expected utility fraction of 0.7111 while $W(0) = \$30 + \$60 = \$90$ induces the fraction 0.8358. Both of these correspond to points on the graph in Figure 5.

Note that having an infusion helps attain goals, even if those goals are at times before the infusion. For example, having no infusion leads to a 4.35% chance of attaining Goal 1 at $t = 5$, and an infusion of an additional \$30 at $t = 0$ boosts this to a 50.51% chance, but if the \$30 infusion happens instead at $t = 10$, the chance of attaining Goal 1 at $t = 5$ is still boosted to 19.06%. This is because the investor is freed up from having to preserve some money from the initial investment to attain later goals.

The later the infusion occurs, the smaller its effect due to the investment having less time to grow. Note, for example, from inspecting the expected utility fractions that an infusion of 60 at $t = 10$ is approximately equal to an infusion of 30 at $t = 0$, and an infusion of 60 at $t = 20$ is approximately equal to an infusion of 30 at $t = 10$. Therefore, the probability of each individual goal generally goes down as the infusion time increases, but the timing of the infusion creates other effects that can counter this. For example, the probability of attaining Goal 6 generally *increases* as the infusion time increases. This makes some intuitive sense since Goal 6, which is at $t = 22$, has a cost of 60, so having an infusion of 60 at the later time of $t = 20$ helps attain this goal more than having the infusion earlier would, when it is more likely to be subject to losses.

6.4 The Effect of Changing the Available Portfolios

Finally, if the investor wishes to increase their ability to fulfill their goals without adding additional money to their portfolio, this can be accomplished by expanding the range of portfolios that are accessible on the efficient frontier by either decreasing μ_{\min} or increasing μ_{\max} . This expanded set of investment possibilities can only increase the expected total utility of the investor. We have chosen the default value of μ_{\min} to be 0.0526, which corresponds to the vertex of the efficient frontier shown in Figure 1. Because it is the vertex, selecting lower values of μ_{\min} make no financial sense, since lower values of μ_{\min} correspond to higher values of σ_{\min} . On the other hand, in Table 12 we can see the effect of increasing the value of μ_{\max} from its default value of 0.0886.

Table 12 shows how having access to more aggressive portfolios increases the expected utility fraction, as it must. It also increases the probability of fulfilling most of the individual goals. The clear exception is Goal 2, whose probability decreases. This decrease happens because the early use of high volatility stocks means the investor is more likely to lose the money they need to safeguard in order to fulfill Goal 2. However, the bigger problem for an investor considering increasing μ_{\max} is the corresponding significant increase in σ_{\max} shown in Table 12, as well as the required significant shorting and going very long for the component positions within such aggressive portfolios.

Table 11: The effect of infusions on the base case shown in Table 3, reprised at the top of the table, for which $W(0) = \$30$. Panel A shows the effect of annual infusions. Panel B shows the effect of a single infusion at the given time $t = t_0$. The infusions' effects on increasing the probabilities of exercising the various goals are shown.

	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
Goal time, t (in years)	5	8	10	11	17	22	24
Goal cost, c (in 1000s of \$)	25	17	15	80	50	60	130
Goal utility, u	1000	2500	500	1500	300	3000	2000

Panel A: Annual infusions

Annual Infusion, $I(t)$	Expected Utility Fraction	Probability of fulfilling goal						
		Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
0	0.5018	0.0435	0.9875	0.2085	0.0308	0.0190	0.7629	0.2311
1	0.6059	0.1468	0.9979	0.4056	0.0709	0.0390	0.9344	0.3895
2	0.6824	0.3419	1.0000	0.6205	0.1080	0.0621	0.9852	0.5409
3	0.7416	0.6278	1.0000	0.7624	0.1494	0.0823	0.9963	0.6312
5	0.8418	0.9811	1.0000	0.9309	0.3386	0.1667	0.9994	0.7946
10	0.9915	1.0000	1.0000	0.9993	0.9996	0.7945	1.0000	0.9854

Panel B: Single Infusion at time $t = t_0$

Infusion time, t_0	Infusion Amount, $I(t_0)$	Expected Utility Fraction	Probability of fulfilling goal						
			Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
0	30	0.7111	0.5051	0.9998	0.6385	0.2354	0.1028	0.9603	0.5458
10	30	0.6351	0.1906	0.9975	0.4950	0.0636	0.0405	0.9724	0.4509
20	30	0.5849	0.1098	0.9902	0.3486	0.0456	0.0258	0.9579	0.3035
0	60	0.8358	0.8646	1.000	0.855	0.5050	0.2233	0.9883	0.7224
10	60	0.7164	0.4459	0.9974	0.9391	0.0930	0.0635	0.9896	0.6006
20	60	0.6388	0.1758	0.9966	0.5229	0.0540	0.0369	0.9986	0.4410

6.5 The Effect of Assigning Utility to Excess Money at $T = 25$

Up until this point, we have assumed that the investor has no use for any excess money left over when the fund is closed at $T = 25$. (That is, $V(W(T)) = U(W) = 0$.) We now consider the effect of valuing this excess money as discussed in Subsection 2.2. In this subsection, we will use the following terminal utility function:

$$V(W(T)) = U(W) = k \left(\frac{1}{1 + be^{-aW}} - \frac{1}{1 + b} \right), \quad (6)$$

where $a, b, k > 0$. Note that $0 < U(W) < k \frac{b}{1+b}$, no matter what the value of a is. As a increases, the investor becomes more risk averse. Generally, we have $a < 0.5$, since behavior where $a > 0.5$ becomes hyper-averse. See Figure 6 for examples of the utility function in equation 6, where $b = 1$ and $k = 2$. Note, as expected, that $U(W)$ is both increasing and concave in W . For the cases in the figure, when $a = 0.01$, $U(W)$ essentially

Table 12: The effect of increasing μ_{\max} — and therefore the available investment portfolios — on the base case shown in Table 3, reprised at the top of the table. The effect of expanding the range of investment portfolios on the probabilities of fulfilling the various goals are shown.

	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
Goal time, t (in years)	5	8	10	11	17	22	24
Goal cost, c (in 1000s of \$)	25	17	15	80	50	60	130
Goal utility, u	1000	2500	500	1500	300	3000	2000

Value for μ_{\max}	Value for σ_{\max}	Expected Utility Fraction	Probability of fulfilling goal						
			Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
0.0886	0.1954	0.5018	0.0435	0.9875	0.2085	0.0308	0.0190	0.7629	0.2311
0.10	0.2551	0.5117	0.0613	0.9756	0.2269	0.0477	0.0238	0.7662	0.2679
0.15	0.5201	0.5462	0.1727	0.9498	0.3394	0.1057	0.0433	0.7665	0.3555
0.20	0.7861	0.6053	0.3747	0.8879	0.4281	0.2231	0.1054	0.8093	0.4670

reaches its maximum value of 1 by $W = 500$, whereas when $a = 0.05$, it is essentially reached by $W = 100$.

We implemented the algorithm with $U(W)$ being given by equation (6) with $a = 0.01$ and $b = 1$, so the maximum utility is $k/2$. Table 13 shows the effect of various values of the magnitude k on $E[W(T)]$, which is the expected final wealth, $E[U(W(T))]$, which is the expected utility from the final wealth that is guaranteed to be less than $k/2$, and the probabilities of attaining the seven goals in the base case from Table 3.

Table 13: The effect of the investor valuing their excess wealth at $T = 25$ on the base case specified in Table 3, reprised at the top of the table. The utility of the excess wealth is based on equation (6), where we fix $b = 1$, $a = 0.01$, and vary the magnitude, k . As k increases, the investor progressively values having excess wealth. The expected excess wealth is $E[W(T)]$. The expected utility of this excess wealth is $E[U(W(T))]$. The utility corresponding to an infinite excess wealth is $k/2$.

	Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
Goal time, t (in years)	5	8	10	11	17	22	24
Goal cost, c (in 1000s of \$)	25	17	15	80	50	60	130
Goal utility, u	1000	2500	500	1500	300	3000	2000

k	$E[W(T)]$	$E[U(W(T))]$	Probability of fulfilling goal						
			Goal 1	Goal 2	Goal 3	Goal 4	Goal 5	Goal 6	Goal 7
0	32.31	0	0.0435	0.9875	0.2085	0.0308	0.0190	0.7629	0.2311
500	36.01	42.70	0.0412	0.9874	0.2058	0.0304	0.0189	0.7608	0.2363
1000	37.70	88.68	0.0390	0.9872	0.2029	0.0300	0.0185	0.7582	0.2413
2000	42.14	193.7	0.0397	0.9861	0.1993	0.0299	0.0152	0.7510	0.2483
5000	61.98	655.9	0.0249	0.9833	0.1614	0.0231	0.0108	0.7248	0.2529
10000	90.35	1779	0.0087	0.9804	0.0918	0.0137	0.0065	0.6785	0.1879




Figure 6: The utility function $U(W) = k \left(\frac{1}{1+be^{-aW}} - \frac{1}{1+b} \right)$, which we can use to value residual wealth, W , at the terminal horizon, T . In the figure, $b = 1$ and $k = 2$, so the maximum value of U is 1. As a increases, the investor becomes more risk averse.

When $k = 500$, the maximum value of $U(W(T))$ is $500/2 = 250$, which is much smaller than the utilities associated to the seven goals, so there is little effect on the probabilities that the goals are attained. Indeed, the probability of attaining Goal 7 is actually increased. This is because the excess money at $T = 25$ is now valued, which pushes the investment portfolios at later times to be more aggressive. These aggressive portfolios lead to more available money, on average, near the end, which is better used for Goal 7, whose utility is 2000, than as excess money, which has a utility of 250 at the most.

As k increases, the expected worth of the excess money increases and the probability of attaining the first six goals decreases since it becomes more worthwhile to have excess funds. The probability for attaining Goal 7, however, continues to increase until the maximum utility from the excess money, $k/2$, exceeds 2000. Once this occurs, the probability of attaining Goal 7 decreases along with the other six goals.