
Efficient Goal Probabilities: A New Frontier

Sanjiv Das¹ Daniel Ostrov¹ Anand Radhakrishnan² Deep Srivastav²

Abstract

In goals-based wealth management (GBWM), an investor looks to maximize the probabilities of attaining each of n goals over time. Because the goals are in competition for potentially limited financial resources, their relative importance must be specified, which we do by assigning utility weights to each goal. Given these weights, dynamic programming can determine both the optimal investment strategy and the optimal strategy for when to fulfill versus forgo each goal. This yields the optimal goal probabilities for fulfilling each goal. By altering the utility weights, we show how to generate the efficient goal probability frontier (EGPF), an $(n - 1)$ -dimensional hypersurface of the optimized goal probability combinations. Just as the classic efficient frontier in mean-variance portfolio optimization allows investors to understand the trade-offs under the best circumstances between their portfolio's mean and variance, the EGPF allows the investor to understand the trade-offs under the best circumstances between the probabilities of attaining each of their goals — *without* needing to see or understand the goals' underlying utility weights. We extend our EGPF framework to determine either the minimum initial wealth or the minimum of a one-parameter family of infusions over time that are needed to attain specified probabilities for each completely or partially fulfilled goal.

Keywords: goals-based wealth management; multiple goals; efficient goal probability frontier; dynamic programming; optimal investment strategy; optimal goals taking strategy; efficient frontier

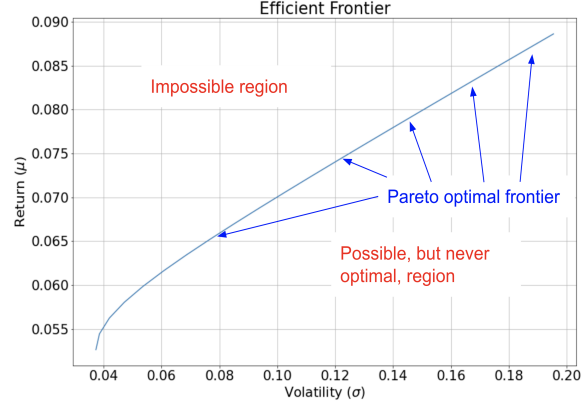


Figure 1. A typical Markowitz (mean-variance) efficient frontier. Points (σ, μ) above the efficient frontier are impossible to obtain. Investment portfolios that lead to points (σ, μ) below the efficient frontier are always suboptimal, so they should never be used. The efficient frontier is a Pareto frontier showing the trade-off between μ and σ under the optimal circumstances. The best location on the efficient frontier at any given time is defined by the investor's preference regarding this trade-off. The specific efficient frontier shown in this figure is used for all the examples in this paper. That is, as we will see, it is the hyperbola from equation (6), where μ and Σ are given in equations (4) and (5).

1. Introduction

The seminal Markowitz (1952) paper defined the notion of the “efficient frontier,” the envelope of best possible portfolios representing the optimal tradeoff of risk and return. The efficient frontier is plotted in (σ, μ) space, where σ and μ are, respectively, the volatility and the expected return of a portfolio. An example of the efficient frontier is shown in Figure 1. Any point, (σ, μ) , above the efficient frontier is impossible to obtain, meaning no portfolios have that (σ, μ) combination. Any point below the efficient frontier is suboptimal, meaning there is a portfolio on the frontier with the same σ , but a higher μ . The frontier itself is Pareto optimal: no location is inherently better than another. The best location on the frontier strictly depends on an investor's personal preferred balance between volatility and expected return.

A vast literature has incorporated and extended mean-variance efficient portfolio optimization, in both static and dynamic settings, and is widely used in retirement plan-

¹Santa Clara University, CA, USA ²Franklin Templeton Inc. Correspondence to: Sanjiv Das <sradas@scu.edu>, Daniel Ostrov <dostrov@scu.edu>, Anand Radhakrishnan <Andy.Radhakrishnan@franklintempleton.com>, Deep Srivastav <deeppratna.srivastav@franklintempleton.com>.

ning (Merton, 1969; 1971; Consiglio, Cocco, and Zenios, 2004; Dempster and Medova, 2011; Koniecz and Mulvey, 2013; Vo and Maurer, 2013; Koniecz, Pisinger, Rasmussen, and Steffensen, 2015; Koniecz, Pisinger, and Weissensteiner, 2016; Brown, Cederburg, and O'Doherty, 2017; Merton and Muralidhar, 2017; Muralidhar, 2018; Simsek, Kim, Kim, and Mulvey, 2018; Martellini, Milhau, and Mulvey, 2019).

More recently, the wealth management industry has advocated the use of goals-based wealth management (GBWM), where investors choose investment portfolios on the Markowitz efficient frontier with an aim to maximize the probability of meeting their financial goals, rather than maximize utility or directly optimize their trade-off between volatility and expected return. Early literature (Browne, 1997; 1999; Shefrin and Statman, 2000; Nevins, 2004; Chhabra, 2005; Das, Markowitz, Scheid, and Statman, 2011; Wang, Suri, Laster, and Almadi, 2011; Das and Statman, 2013) laid the groundwork for GBWM. A wide ranging book by Brunel (2015) offers a framework. In GBWM, the notion of risk, usually deemed to be the volatility of the portfolio return, is instead interpreted as the chance that an investor does not meet their goal(s). The two views may conflict; that is, decreasing risk in the sense of lowering volatility may increase risk in the sense of an investor not reaching their goals. This happens, for example, for young, poor investors who put all of their retirement assets into cash.

Goals-based wealth management implementations vary, and several recent works have offered theoretically founded, yet practical approaches (Deguest, Martellini, Milhau, Suri, and Wang, 2015; Alexander, Baptista, and Yan, 2017; Das, Ostrov, Radhakrishnan, and Srivastav, 2018; 2020; Das and Ross, 2020; Kim, Kwon, Lee, Kim, and Lin, 2020; Martellini, Milhau, and Mulvey, 2020; Parker, 2020; Das, Ostrov, Casanova, Radhakrishnan, and Srivastav, 2021; Parker, 2021).

In Das, Ostrov, Radhakrishnan, and Srivastav (2021), the authors present an algorithm for dynamic optimization over long horizons with multiple goals, where each goal can be fulfilled at a specified time and cost. When there is sufficient money to meet all lifetime goals, the GBWM problem is trivial and the probability of meeting each goal is essentially equal to one. When the wealth in the portfolio is insufficient to meet all goals, however, the investor will need a way to weight their goals with utilities to instruct the algorithm which goals to prioritize by maximizing the expected utility of fulfilled goals over time. For instance, in the presence of two goals with relative utility weights of, say, 3:1, the algorithm will return a higher probability of achieving the first goal than the second compared to if the relative utility weights were, say, 1:3. Of course, either way the probabilities of fulfilling the two goals will be maximized while respecting these utility weights. Since different ratios of utility weights result in different pairs of maximized proba-

bilities, we can plot all the possible maximized probability pairs as we vary the underlying ratio of the utility weights from 0 to ∞ . The plot of all these maximized probability pairs forms the “efficient goal probability frontier” (EGPF), which is a one-dimensional curve in the two-dimensional space formed by the probabilities of fulfilling each of the two goals.

We note the following properties of the efficient goal probability frontier (EGPF):

1. For two goals, the EGPF is a curve. For n goals, the EGPF is an $(n - 1)$ -dimensional hypersurface.
2. Each point on the EGPF is mapped from a predetermined set of utility weights that reflect the relative importance of each goal to the investor.
3. Behind each point of the EGPF are two optimized strategies: 1) an optimized investment strategy and 2) an optimized goals-taking strategy. The optimized investment strategy dynamically, optimally moves the investor among a given set of possible investment portfolios as time and the investor's wealth change. It uses mean-variance efficient investment portfolios when possible, as in the cases presented in this paper, but can just as easily work with other sets of investment portfolios when mean-variance efficient investment portfolios are not available. The optimized goals-taking strategy best determines which goals to fulfill and which to forgo so that more important goals may be fulfilled later.

As mentioned earlier, the Markowitz mean-variance efficient frontier (EF) is a Pareto optimal frontier where (σ, μ) points above the EF are impossible to obtain, and points below the EF are possible, but suboptimal, to obtain. The EGPF has the same nature. For n goals, the EGPF is an $(n - 1)$ -dimensional Pareto optimal frontier where goal probabilities (p_1, p_2, \dots, p_n) that correspond to points lying above the EGPF are impossible to obtain, while those below the EGPF are possible but suboptimal. That is, a suboptimal goals-taking strategy and/or a suboptimal investment strategy will lead to a point below the EGPF. Just as an investor must choose which point on the EF (that is, which investment portfolio) best fits their preferences, they must also decide which point on the EGPF (specifying optimal combinations of probabilities for attaining each of their n goals) best fits their preferences.

What is key here is that the EGPF is stated strictly in terms of probabilities, not in terms of utilities, so the investor is *able* to specify their preferences. That is, most investors understand the idea of having to make trade-offs between probabilities intuitively, but almost no investors know how to assign utility values to each of their goals. Because the EGPF only uses utilities behind the scenes, investors no longer need to think about or understand utilities. Instead, they can focus strictly on their goal probabilities, allowing

them to knowledgeably select their best option.

Further, for a fixed initial investment, we can find a point on the EGPF, which we call the proximity point, that is close to an investor's desired goal probability point when the desired goal probability point lies above the EGPF. This helps optimize the use of the initial investment given the investor's desired goal preferences.

If the investor can devote more money to the initial investment, we can instead find the minimum initial investment needed to obtain the desired goal probability point. This enables the investor to neither overspend nor underspend to obtain their desired goal probabilities. Geometrically, it corresponds to finding the EGPF that contains the desired goal probability point, since the EGPF moves upwards towards higher probabilities as the initial investment increases. Alternatively, the investor may not be able to access additional money initially, but may instead be able to contribute money over time. In this case, the method to determine the minimum additional initial investment can easily be altered to instead determine the minimum of any given one-parameter family of infusions over time where the desired goal probability point is obtained.

While the EGPF can be easily visualized for two goals (in 2-D) or for three goals (in 3-D), it cannot be easily visualized in higher dimensions. However, our algorithms for finding the proximity point and for finding the minimum initial wealth or minimum infusions to obtain desired goal probabilities works for n goals, where n can be far larger than 2 or 3. Further, this can be extended to accommodate the important case of partial goals. For example, a couple may have a goal of spending \$40,000 on a car seven years from now, but also be open to a partial goal where they spend \$35,000 on a less nice car or another partial goal where they only spend \$28,000 on the car. The partial goals will bring less happiness to the investor, but the money may be better off being diverted to other concurrent goals or saved to give a higher likelihood of attaining more important goals at later times.

The paper proceeds as follows. Section 2 discusses how, in the context of a single goal and static one-period optimization, we can generate goal probability level curves (GPLCs), which then determine the single point that comprises the EGPF in this simple context. Section 3 extends this considerably to multiple goals with dynamic (multi-period) optimization. Section 4 shows how to determine the proximity point on an EGPF, as well as the minimal initial investment or minimal infusions needed to obtain desired goal probabilities for full or partial goals. Further, it explores the sensitivity of these results to prohibiting short positions or to increasing either the expected returns or the covariance matrix of the investments. Section 5 contains a closing discussion.

2. Goal Probability Level Curves For Static Optimization With A Single Goal

Das, Ostrov, Radhakrishnan, and Srivastav (2018) derived a mathematical relationship between the Markowitz efficient frontier (EF) and goals-based investing for a given single goal. The single goal was to attain enough portfolio growth over a set time period to enable fulfilling (that is, paying for) a desired financial goal at end of the time period. They showed how to locate the static investment portfolio on the efficient frontier that maximizes the probability of achieving the given goal. This investment portfolio sits at a common point of tangency between the hyperbolic, concave curve in the (σ, μ) plane for the EF and the convex curve of all (σ, μ) combinations that correspond to the maximum attainable probability of fulfilling the given goal. This latter, convex curve is a specific example of a goal probability level curve (GPLC), which is defined as any iso-probability curve in the (σ, μ) plane.

We further explain the nature of a GPLC, following Das, Ostrov, Radhakrishnan, and Srivastav (2018). We assume a wealth (W) growth process that follows geometric Brownian motion:

$$\frac{dW}{W} = \mu dt + \sigma dB, \quad (1)$$

where μ is the continuous mean growth rate, σ is the volatility, and dB is the standard Brownian motion differential. The solution to this stochastic differential equation in (1) is

$$W(t) = W(0) \exp \left[\left(\mu - \frac{1}{2}\sigma^2 \right) t + \sigma\sqrt{t} Z \right], \quad (2)$$

where $Z \sim N(0, 1)$ is a standard normal random variate.

Rearranging equation (2) so that μ is isolated on the left-hand side and replacing Z with $-z_0$ yields

$$\mu = \frac{1}{2}\sigma^2 + \frac{z_0}{\sqrt{t}}\sigma + g, \quad (3)$$

where $g = \frac{1}{t} \ln \left(\frac{W(t)}{W(0)} \right)$ is the constant, continuous growth rate per annum required for the initial investment $W(0)$ to grow to the target goal $W(t)$, and z_0 is the unique value such that $Pr(Z \leq z_0)$ is equal to the "goal probability," which is the probability that $W(0)$ will grow to at least $W(t)$, the target goal. If we fix the values of g , t , and z_0 , then equation (3) defines a convex parabolic relationship for μ as a function of σ . This convex parabola in the (μ, σ) plane is denoted as the goal probability level curve (GPLC).

As an example, assume we have access to three index funds whose expected returns are given by the vector

$$\mu = \begin{bmatrix} 0.0493 \\ 0.0770 \\ 0.0886 \end{bmatrix} \quad (4)$$

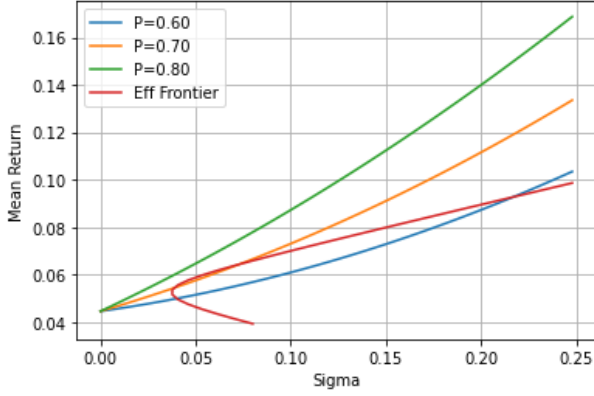


Figure 2. The efficient frontier from Figure 1 along with three goal probability level curves (GPLCs). The initial wealth is $W(0) = \$100,000$. The target wealth after five years is $W(5) = \$125,000$. The three GPLCs shown correspond to goal probabilities of 60%, 70%, and 80%, meaning, for example, any investment portfolio whose (μ, σ) pair is on the 60% GPLC has a 60% chance that $W(5) \geq \$125,000$.

and whose return covariance matrix is

$$\Sigma = \begin{bmatrix} 0.0017 & -0.0017 & -0.0021 \\ -0.0017 & 0.0396 & 0.0309 \\ -0.0021 & 0.0309 & 0.0392 \end{bmatrix}. \quad (5)$$

The EF is the hyperbola

$$\sigma = \sqrt{a\mu^2 + b\mu + c}, \quad (6)$$

with $a = \mathbf{h}^\top \Sigma \mathbf{h}$, $b = 2\mathbf{g}^\top \Sigma \mathbf{h}$, and $c = \mathbf{g}^\top \Sigma \mathbf{g}$, where $\mathbf{g} = \frac{l\Sigma^{-1}\mathbf{1} - k\Sigma^{-1}\boldsymbol{\mu}}{lm - k^2}$, $\mathbf{h} = \frac{m\Sigma^{-1}\boldsymbol{\mu} - k\Sigma^{-1}\mathbf{1}}{lm - k^2}$, and $\mathbf{1}$ is a unit vector of ones, and where the scalars k , l , and m are defined by $k = \boldsymbol{\mu}^\top \Sigma^{-1} \mathbf{1}$, $l = \boldsymbol{\mu}^\top \Sigma^{-1} \boldsymbol{\mu}$, and $m = \mathbf{1}^\top \Sigma^{-1} \mathbf{1}$. Figure 1 shows the efficient frontier for the $\boldsymbol{\mu}$ and Σ given in equations (4) and (5).

Now assume that the initial wealth is $W(0) = \$100,000$ and the target wealth after five years is $W(5) = \$125,000$. Figure 2 shows the efficient frontier from Figure 1 along with the three different GPLCs corresponding to goal probabilities of 60%, 70%, and 80%, which are created by increasing the value of z_0 . We see that we can do much better than a 60% goal probability, since so much of the EF lies above the 60% GPLC. Only a little of the EF lies above the 70% GPLC, however, and no attainable investment portfolio can achieve a goal probability of 80%, since the 80% GPLC lies strictly above the EF.

It is therefore clear that the highest possible goal probability corresponds to a GPLC that shares a single point in common with the EF. Further, this point is a common point of tangency for the EF and this GPLC, as shown in Figure 3, and more importantly, it defines the μ and σ of the optimal static

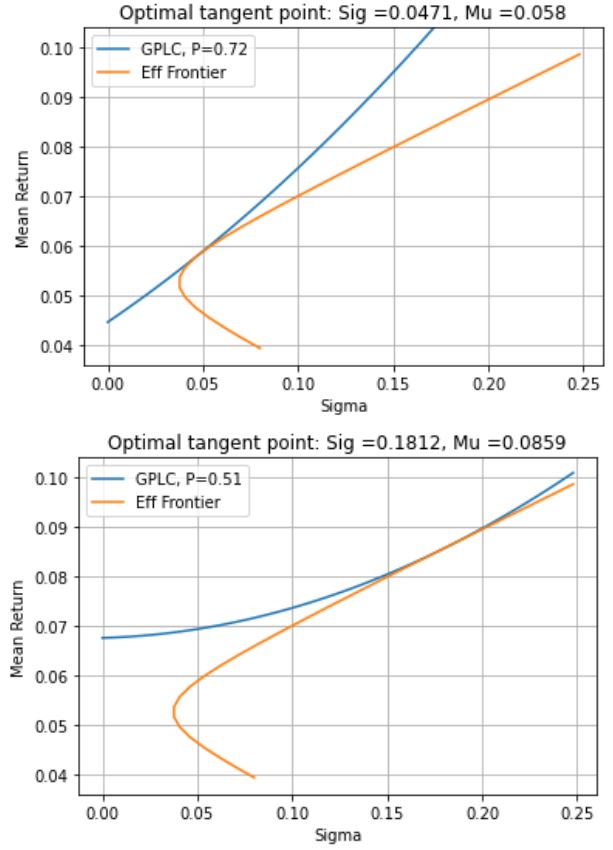


Figure 3. The optimal goal probability investment portfolio. This resides at the common point of tangency between the goal probability level curves (GPLCs) and the efficient frontier. For the top panel, the initial wealth is $W(0) = \$100,000$, the target wealth after five years is $W(5) = \$125,000$, and the optimal goal probability is 72%. For the bottom panel, we have the same parameters, including the initial wealth, but the target wealth, which is now after six years, is $W(6) = \$150,000$. In this case the optimal goal probability becomes 51%, and the optimal portfolio is more aggressive than in the top panel.

investment portfolio. We know this optimal GPLC corresponds to a goal probability just above 70%, but how do we determine the exact value of the optimal goal probability and the exact optimal point on the EF?

Based on [Das, Ostrov, Radhakrishnan, and Srivastav \(2018\)](#), we re-arrange equation (3) to isolate z_0 ,

$$z_0 = \frac{1}{\sigma} \left(\mu - \frac{1}{2} \sigma^2 - g \right) \sqrt{t}, \quad (7)$$

and then look to maximize z_0 , while constrained to remain on the EF (equation 6). This constrained optimization problem yields a Lagrange Multipliers formulation that results in a cubic polynomial equation for the optimal value of μ :

$$c_3 \mu^3 + c_2 \mu^2 + c_1 \mu + c_0 = 0 \quad (8)$$

where $c_3 = a^2$, $c_2 = \frac{3ab}{2}$, $c_1 = ac + \frac{b^2}{2} - b - 2ag$, and $c_0 = \frac{bc}{2} - 2c - bg$. This cubic polynomial has three potential roots for μ , but it is guaranteed that one, and only one, of these roots will be above the EF vertex's μ value. The point on the upper envelope of the EF with this root's μ value is the optimal static portfolio.

For our example, this optimal solution point is shown in the top panel of Figure 3, where we see the optimized goal probability is 72%, which is just above 70%, as expected. This 72% goal probability is obtained by using the optimal investment portfolio on the EF, which has an expected return of 5.80% and a volatility of 4.71%. The lower panel in Figure 3 is for the same circumstances, but a different, more difficult target goal: instead of five years, we have six years, but we look to grow our initial \$100,000 to \$150,000 instead of just \$125,000. Given the more difficult goal, we expect both a lower optimal goal probability and a more aggressive optimal investment portfolio. The lower panel confirms and quantifies this: the optimal goal probability is only 51%, and the optimal investment portfolio (which is on the same EF as before) has an expected return of 8.59% and a volatility of 18.12%.

The optimal solution point's GPLC probability corresponds to the single point that comprises the entire efficient goal probability frontier (EGPF) in this section's context, as we now explain: Because the number of goals, n , equals one here, we have a one-dimensional probability space, which is simply a line segment from 0 to 1. The EGPF is an $(n - 1)$ -dimensional hyperplane in general, meaning a zero-dimensional hyperplane in our current context, which is a point. That point is the location on the line segment that corresponds to the value of the optimal probability of fulfilling the goal, which is the optimal solution point's GPLC probability.

As we will see later, each point on the EGPF comes with two optimal strategies. The first is the optimal investment strategy. In our current context, this is the optimal static portfolio, which corresponds to the optimal solution point's location on the EF. The second is the optimal goals taking strategy. In our current context, this is obvious: if the investor has enough money to fulfill their goal, they should fulfill it, since, when there is only one goal as is the case here, there is no later goal that is more likely to be fulfilled if we forego our single goal.

3. Efficient Goal Probability Frontier With Dynamic Optimization And Multiple Goals

The analysis in the previous section presented a solution to maximize the probability of reaching a single goal in a static investment portfolio framework. In this section, we extend this framework considerably, both to encompass multiple goals instead of one goal, and to use dynamic optimization over multiple time periods instead of static optimization.

Dynamic optimization in the context of the remainder of this paper means: (i) Investment optimization, where the investment portfolio can change with each time period and the investment optimization strategy must take into account the fact that the investment portfolio can be changed with each future time period. (ii) Goals-taking optimization, where each goal at each time period must be optimally chosen to be fulfilled or unfulfilled so as to optimize the expected value of the sum of the utilities of the fulfilled goals over both current and future time periods.

Because we now consider multiple goals, we must weigh the relative importance of these goals to the investor, which we do by assigning a utility value to each goal. Each set of relative utility weights for the goals will correspond to a point on the efficient goal probability frontier (EGPF), which represents the set of optimal probabilities for the goals given these relative utility weights. In this section we explore the geometric nature of the EGPF in the case of two goals and three goals.

3.1. An Example With Two Goals

Consider an investor with an initial wealth of \$50,000 and two goals: Goal 1 is to spend \$100,000 for a luxury car 10 years from now. Goal 2 is to spend \$150,000 for their child's college 20 years from now. With two goals, the optimized goal probabilities depend on the relative importance (utilities) that the investor assigns to each goal. If we assign a utility of $U_1 = 100$ to the car goal and a utility of, say, $U_2 = 20$ to the college goal, we can use the dynamic programming algorithm described in Appendix A to maximize the expected attained utility, which will generate the optimal goal probabilities for the car and college goals. If we keep $U_1 = 100$ but increase U_2 to, say, 80, the algorithm will return a lower optimized goal probability for the car goal and a higher one for the college goal.

Indeed, as the ratio $\frac{U_2}{U_1}$ is increased from 0 to infinity, we expect to see the corresponding optimized goal probabilities shift from favoring Goal 1 to favoring Goal 2. Specifically, the optimized goal probabilities (p_1, p_2) will trace a curve in the plane whose axes are the Goal 1 probability and the Goal 2 probability. This curve forms the efficient goal probability frontier (EGPF) shown in Figure 4.

Note that even though assigned utility values underlie our analysis, the EGPF is strictly in terms of probabilities. This is key because investors do not have a clear sense of what utility values to assign to goals, but they can understand the trade-offs between the optimal probability of fulfilling one goal versus another. The EGPF makes those trade-offs explicit, allowing an investor to make an informed choice that conforms to their preferences.

There are clear parallels between the EGPF in Figure 4 and the Markowitz efficient frontier (EF) in Figure 1. Both frontiers are Pareto optimal, so no point on a frontier is in-

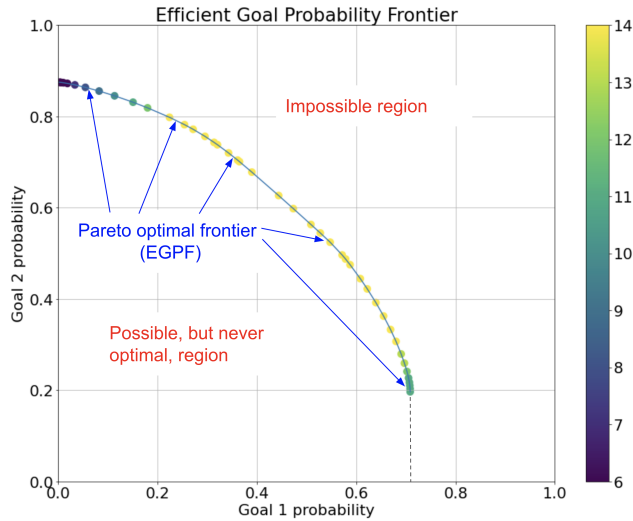


Figure 4. The efficient goal probability frontier (EGPF) for two goals. This EGPF is for an investor starting with \$50,000 with access to 15 investment portfolios discussed in the text. The EGPF curve shows various possible pairs of dynamically optimized goal probabilities for two goals, (i) buying a luxury car in 10 years time at a cost of \$100,000 and (ii) saving for college so as to have \$150,000 in 20 years time. That is, the EGPF shows the trade-off between the probability of fulfilling Goal 1 versus Goal 2 under the optimal investment and goals-taking strategies. The colors on the EGPF correspond to which of the 15 investment portfolios are initially optimal. The EGPF, combined with the dashed vertical line when the Goal 2 probability is below approximately 20%, partitions the graph into a region that is impossible to attain and a region that is possible but suboptimal. Note the heavy parallel between this figure and Figure 1, which described the Markowitz efficient frontier (EF).

herently better or worse than another. That is, the best point on either frontier is a matter of investor-specific preferences. It is impossible to attain a point above either frontier. It is possible to attain a point below either frontier, but never desirable, as these points represent inherently suboptimal investment and/or goals-taking decisions. Of course, the EF is always a 1-D curve residing in the 2-D space formed by μ and σ , whereas for n goals, the EGPF is an $(n - 1)$ -dimensional hypersurface residing in the n -dimensional space formed by the goal probabilities of each goal.

Unlike the static analysis in Section 2 where we could select any single point on the EF, our analysis here allows us in each time period to dynamically jump between a fixed number of investment portfolios on the EF in Figure 1. To obtain the EGPF in Figure 4, the dynamic programming algorithm had access to 15 investment portfolios. These portfolios were determined by equally spacing 15 values for μ from 5.26%, which corresponds to the vertex of the EF hyperbola, to 8.86%, the highest expected value of the three index funds from equation (4), and then selecting the 15 investment portfolios (numbered portfolio 0 to portfolio 14) on the EF in Figure 1 with these μ values.

While the optimal choice among these 15 investment portfolios varies with time due to varying wealth, we are able to show the initial optimal investment in Figure 4, since we know the initial wealth is \$50,000. In the figure we see that when $\frac{U_1}{U_2}$ is near zero, meaning Goal 2 is far more important than Goal 1, it is optimal to initially select portfolio 6, a relatively conservative portfolio, since we only worry about obtaining Goal 2 for \$150,000 after 20 years, not Goal 1. As $\frac{U_1}{U_2}$ increases to values closer to 1, we see the initial portfolio increase to portfolio 14, the most aggressive portfolio, since we have a stronger reason to accept more volatility in the hopes of attaining both goals. As $\frac{U_1}{U_2}$ increases towards infinity, we see the initial portfolio come back down to portfolio 11. This corresponds to only worrying about obtaining Goal 1, which means obtaining \$100,000 after 10 years.

Because Goal 1 is harder to reach than Goal 2, we optimally require a more aggressive initial portfolio (number 11) than portfolio 6, the optimal initial portfolio when we only worried about obtaining Goal 2. We can also see this in Figure 4, where the optimal probability of attaining Goal 2 is 87.5% (on the top, left point of the EGPF), while the optimal probability of attaining Goal 1 is 70.9% (on the bottom, right point of the EGPF). At the top, left point, the optimal goals-taking strategy is to never fulfill Goal 1, so as to maximize the probability of being able to fulfill Goal 2, so we see that the optimal probability of attaining Goal 1 at this point is zero. In contrast, at the bottom, right point, the optimal goals-taking strategy is to fulfill Goal 1 at time 10 if the money exists, but then use any excess money to try to obtain Goal 2, which happens 19.7% of the time. In this case, there is no advantage towards fulfilling Goal 1 by using a strategy that further reduces the chance of fulfill-

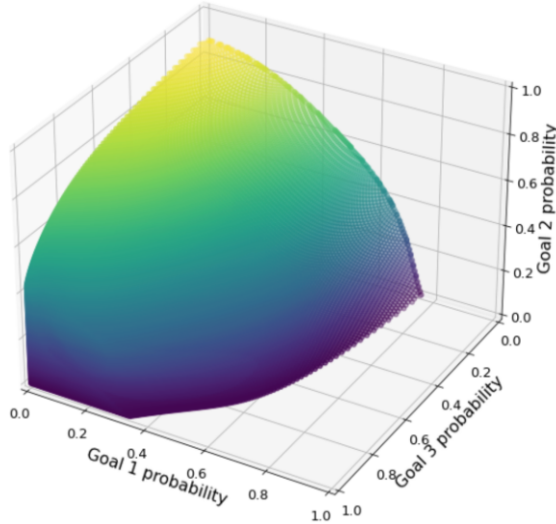


Figure 5. Efficient goal probability frontier for three goals. The investor starts with \$50,000. The surface shows various possible combinations of dynamically optimized goal probabilities over three goals, (i) buying a luxury car in 10 years time at a cost of \$100,000, (ii) saving for college so as to have \$150,000 in 20 years time, and (iii) taking an extended vacation in 15 years time at a cost of \$70,000.

ing Goal 2. This is indicated by the dashed, vertical line in Figure 4 that extends from the bottom, right point. The dashed line (along with the EGPF) demarcates the region that is possible, but never optimal, from the region that is impossible to attain.

3.2. Extension To Three Goals

We now examine a case with three goals. In Subsection 3.1, we considered Goal 1, a car goal costing \$100,000 in 10 years time, and Goal 2, a college goal costing \$150,000 in 20 years time. To this we add Goal 3, an extended vacation goal costing \$70,000 in 15 years time. Using the approach in Appendix A, the optimized probabilities of fulfilling the three goals are obtained for different utility weights assigned to each goal. This traces out a 2-D EGPF surface in 3-D probability space, as shown in Figure 5. The EGPF surface gives the best possible combinations of goal probabilities across the three goals. For example, an investor mostly interested in making sure they fulfill Goal 2 and less concerned about Goals 1 and 3 might choose the dynamic strategy associated with a yellow point near the top of the surface. Note that when we set the probability of obtaining Goal 3 to zero, the cross-section of the corresponding plane with the EGPF surface is the 1-D EGPF curve in Figure 4.

4. Working With Multiple Goals And Desired Probabilities For Each Of These Goals

The EGPF (efficient goal probability frontier) makes it clear to investors what the trade-offs are between their goals under optimal circumstances. Once the investor selects a point on the EGPF that best satisfies their personal preference within these trade-off constraints, we can note the corresponding underlying utility values assigned to that point for each goal and, using Appendix A, we can determine both the corresponding optimal investment strategy and the optimal goals-taking strategy for the investor.

But what happens if we have four or more goals? In this case, the EGPF becomes hard to directly visualize, since it exists in four or higher-dimensional space. The ideas behind the EGPF, however, can still be extended to determine how to best satisfy an investor's needs, as we show in this Section. Further, our approach will extend to the case of partial goals, which are acceptable alternatives to full goals for the investor that have the advantage of costing less but the disadvantage of providing less utility should the partial goal be fulfilled instead of the full goal.

4.1. Generating Optimal Probabilities For n Goals When The Initial Wealth Is Fixed

The investor can be asked to give desired probabilities for attaining each of their n goals. This corresponds to a single point in n dimensional probability space, which we will call the “desired point.” We look to determine a point on the EGPF that is “near” (in a sense we will define) to the desired point, even if the geometry cannot be visualized easily when $n \geq 4$. We will call this point on the EGPF the “proximity point” and the algorithm that generates it the “proximity algorithm,” which we describe in this subsection.

We start with some notation before describing the proximity algorithm. We denote the desired point by \mathbf{d} , an n -vector whose components, d_i , are the desired probabilities of attaining goal i , where $i = 1, 2, \dots, n$. We will use the n -vector \mathbf{U} to denote the utilities assigned to each of the n goals, and we will use the n -vector \mathbf{p} to denote a point on the EGPF for a given initial investment $W(0)$, meaning \mathbf{p} represents an attainable, Pareto optimal set of probabilities for each of the n goals, given an initial investment of $W(0)$. Finally, we define the “distance” between two points, such as \mathbf{d} and a point \mathbf{p} , in the normal sense; that is, the distance is $\|\mathbf{d} - \mathbf{p}\|$,

where $\|\mathbf{v}\|$, the norm of an n -vector, is defined by $\sqrt{\sum_{i=1}^n v_i^2}$.

To determine \mathbf{p}_{prox} , the proximity point on the EGPF to the desired point \mathbf{d} , we employ the following iterative algorithm: We start with any initial \mathbf{U} .¹ We then compute the point

¹While, for simplicity, we generally set all the initial $U_i = 100$, we also found that in practice, regardless of the initial \mathbf{U} chosen, the number of iterations, and therefore the computational time for

\mathbf{p} on the EGPF that corresponds to the initial \mathbf{U} by using the algorithm given in Appendix A. After that, we use the following iteration formula (where we later define ΔU) to determine a new value for \mathbf{U} :

$$\mathbf{U}_{\text{new}} = \mathbf{U}_{\text{old}} + \Delta U(\mathbf{d} - \mathbf{p}_{\text{old}}). \quad (9)$$

We then return to using Appendix A to determine the \mathbf{p}_{new} that corresponds to \mathbf{U}_{new} . Setting \mathbf{U}_{new} and \mathbf{p}_{new} to now be \mathbf{U}_{old} and \mathbf{p}_{old} , we can start a new iteration using equation (9).

The iteration formula in equation (9) essentially uses gradient descent in \mathbf{p} to evolve \mathbf{U} . While this is an unusual approach, it proves quite practical in our context. Not only is it extremely quick to compute, it has the key property that for each of the n components where $p_{\text{old},i} < d_i$, we have that $U_{\text{new},i} > U_{\text{old},i}$, and where $p_{\text{old},i} > d_i$, we have that $U_{\text{new},i} < U_{\text{old},i}$. This means, as a practical matter, that as the algorithm evolves, eventually either all the components of \mathbf{p} become greater than or equal to their counterpart components in \mathbf{d} or all the components of \mathbf{p} become less than or equal to their counterpart components in \mathbf{d} .

Initially, we set ΔU to just below the average value of the initial \mathbf{U} components. Two things can go wrong when we apply equation (9). The first is that some of the computed components of \mathbf{U}_{new} may be negative, in which case, we simply reset those components to take small positive values. The second is that ΔU may be too large, which we define as resulting in the distance between \mathbf{d} and \mathbf{p}_{new} being larger than the distance between \mathbf{d} and \mathbf{p}_{old} . When this happens, the algorithm shrinks ΔU to 70% of its previous value and then recomputes \mathbf{U}_{new} from \mathbf{U}_{old} using equation (9) with the new ΔU value. This is repeated as many times as is necessary to ensure that the distance between \mathbf{d} and \mathbf{p}_{new} never grows between iterations.

When do we stop our proximity algorithm? It is a mistake to require that the distance between \mathbf{U}_{new} and \mathbf{U}_{old} becomes sufficiently small, because if $\mathbf{U}_{\text{new}} = k\mathbf{U}_{\text{old}}$, where k is any positive constant, both \mathbf{U}_{new} and \mathbf{U}_{old} will result in the same \mathbf{p} . This is because only the ratios of the components in \mathbf{U} to each other matter, not their actual values. We therefore want to stop once \mathbf{U}_{new} becomes sufficiently close to being a constant multiple of \mathbf{U}_{old} , which, from equation (9), corresponds to the case where

$$\mathbf{U}_{\text{new}} = k\mathbf{U}_{\text{old}} = \mathbf{U}_{\text{old}} + \Delta U(\mathbf{d} - \mathbf{p}_{\text{old}}), \quad (10)$$

or, after rearrangement,

$$\alpha\mathbf{U}_{\text{old}} = \mathbf{d} - \mathbf{p}_{\text{old}}, \quad (11)$$

where $\alpha = \frac{k-1}{\Delta U}$. To remove the unknown value of α , we take the norm of both sides of equation (11) and then divide equation (11) by it. This cancels the positive α factor,

the proximity algorithm to converge, stayed approximately the same.

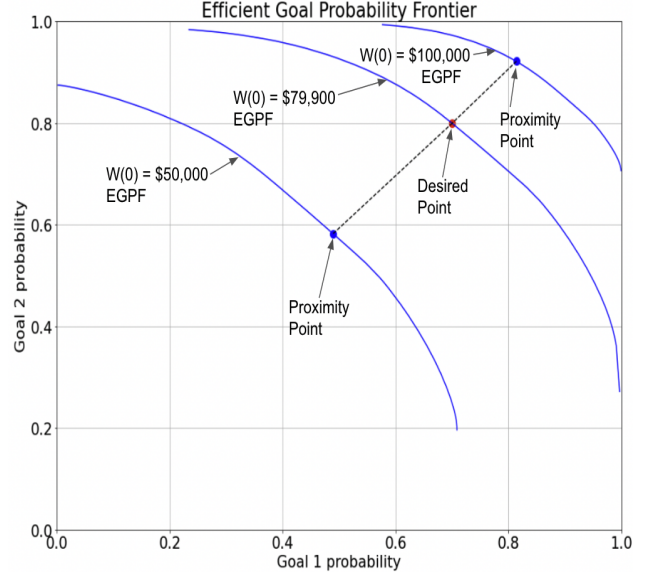


Figure 6. We consider the same two goals and available investments that were used in Subsection 3.1 and Figure 4. The figure shows three efficient goal probability frontiers (EGPFs) corresponding to what is possible if the investor starts with $W(0) = \$50,000$, $\$79,900$, and $\$100,000$. The desired probability point (in red) corresponds to a desired 70% chance of obtaining Goal 1 and a desired 80% chance of obtaining Goal 2. The proximity points from successive EGPFs, shown here just for the $W(0) = \$50,000$ and $\$100,000$ EGPFs, can be used to determine that $W(0) = \$79,900$ is the minimal initial investment needed to obtain the desired probabilities.

leaving

$$\frac{\|\mathbf{U}_{\text{old}}\|}{\|\mathbf{U}_{\text{old}}\|} = \frac{\|\mathbf{d} - \mathbf{p}_{\text{old}}\|}{\|\mathbf{d} - \mathbf{p}_{\text{old}}\|}.$$

It follows that we stop our proximity algorithm when the metric

$$\left\| \frac{\|\mathbf{U}_{\text{old}}\|}{\|\mathbf{U}_{\text{old}}\|} - \frac{\|\mathbf{d} - \mathbf{p}_{\text{old}}\|}{\|\mathbf{d} - \mathbf{p}_{\text{old}}\|} \right\|$$

becomes sufficiently small, and when this occurs we have our proximity point $\mathbf{p}_{\text{prox}} = \mathbf{p}_{\text{old}}$.

The proximity point may give the investor an option that satisfies them, but if they believe one or more of its goal probabilities (that is, components) are too small (or too big), they can change the desired point, \mathbf{d} , to better align with their desires, knowing that each increase to a single goal's probability will generally decrease the other goals' probabilities once the new corresponding proximity point is determined from rerunning our proximity algorithm.

Although the proximity algorithm works for n goals, we can visualize its effect by looking at our two goal example from Subsection 3.1. Recall that the EGPF for this example where $W(0) = \$50,000$ is given in Figure 4. This same EGPF is

reprinted in Figure 6 as the $W(0) = \$50,000$ EGPF. In this figure, we also see the desired point $\mathbf{d} = (0.7, 0.8)$; that is, the investor is looking for at least a 70% chance of attaining Goal 1, which is a \$100,000 luxury car after 10 years, and at least an 80% chance of attaining Goal 2, which is \$150,000 for their child's college in 20 years. The proximity point, \mathbf{p}_{prox} , determined from our algorithm, which is shown in Figure 6, is $(0.4894, 0.5828)$. This proximity point is a distance of 0.3025 from the investor's desired point. The investor can now do one of three things:

1. They can accept having a 48.94% chance of fulfilling Goal 1 and a 58.28% chance of fulfilling Goal 2, in which case, the \mathbf{U} that was used to generate \mathbf{p}_{prox} also generates the optimal investment strategy and optimal goals-taking strategy that are needed to obtain these probabilities, as described in Appendix A.
2. They may decide that they are willing to accept a reduction in the probability of fulfilling one of their goals so that they can increase the probability of fulfilling the other goal. In this case, they will rerun the proximity algorithm with a new desired point \mathbf{d} , which will produce a new \mathbf{p}_{prox} on the $W(0) = \$50,000$ EGPF.
3. They may have more money available, in which case they may be able to obtain their desired probabilities of 70% for Goal 1 and 80% for Goal 2 by adding this money to create a larger $W(0)$. But how much additional money is necessary? The next subsection shows how we can use the proximity algorithm to determine this.

4.2. Determining The Minimum Initial Wealth To Reach The Desired Probabilities

In this subsection, we show how to determine the minimum initial wealth needed to obtain desired probabilities, \mathbf{d} , for fulfilling the investor's goals by applying the secant method. For the secant method, we use the initial wealth, $W(0)$, as the input and the signed distance from \mathbf{d} of the proximity point \mathbf{p}_{prox} (determined from the proximity algorithm in the previous subsection) as the output. Repeated iteration of the secant method will produce proximity points that converge to \mathbf{d} (that is, the unique root of the signed distance) and $W(0)$ values that converge to the minimum initial wealth.

More specifically, as discussed in the previous subsection, the proximity algorithm will produce a proximity point, \mathbf{p}_{prox} , whose components are all less than their corresponding components in \mathbf{d} , in which case the portfolio is underfunded and cannot obtain \mathbf{d} , or all the components in \mathbf{p}_{prox} are greater than (or equal to) their corresponding components in \mathbf{d} , in which case the portfolio is overfunded (or funded). If the portfolio is overfunded, we define the signed distance to be the negative of the distance between \mathbf{p}_{prox} and \mathbf{d} ; if it is underfunded, the signed distance is the (positive) distance.

We run the proximity algorithm using an initial $W(0)$. If that leads to a positive signed distance, meaning we are underfunded, we rerun the proximity algorithm with twice the initial $W(0)$ used previously; otherwise, if we are overfunded, we rerun the proximity algorithm with half the initial $W(0)$ used previously. This gives us two points for which the input is (the initial) $W(0)$ and the output is the corresponding signed distance between \mathbf{p}_{prox} and \mathbf{d} .

From these two points, we can iterate using the secant method to converge to the $W(0)$ that corresponds to the root of the signed distance. In the secant method, we plot the last two calculated points on the $(W(0), \text{signed distance})$ plane, draw the line that goes through the two points (that is, the secant line), find the $W(0)$ value on this line that corresponds to a signed distance of 0, and use that $W(0)$ for the next point, where the signed distance corresponding to this $W(0)$ is determined from the proximity algorithm.

As an example, we continue the two goals case discussed in the previous subsection and shown in Figure 6. Recall that when $W(0) = \$50,000$, the proximity algorithm determined the corresponding signed distance was 0.3025. Since we have underfunding, we next choose $W(0) = \$100,000$. Using this $W(0)$, the proximity algorithm determines that $\mathbf{p}_{\text{prox}} = (0.814, 0.922)$, which corresponds to a signed distance of -0.1670, which is negative because we now have overfunding. In Figure 6 we have calculated the entire $W(0) = \$100,000$ EGPF strictly for visualization purposes. We note, however, that the proximity algorithm does not require the entire EGPF. The proximity algorithm just calculates a few points on the EGPF as it converges to the proximity point, which is why it is so effective with n goals, even if n is large.

We next apply the secant method to the two points $(\$50,000, 0.3025)$ and $(\$100,000, -0.1670)$. The root of the secant line connecting these two points is determined from

$$\frac{W(0) - 50,000}{0 - 0.3025} = \frac{100,000 - 50,000}{-0.1670 - 0.3025},$$

which gives $W(0) = \$82,220$. Applying the proximity algorithm to $W(0) = \$82,220$ gives a signed distance of -0.0338, which is much closer to 0 than our previous two points. We then rerun the secant method with our two most recent points, $(\$100,000, -0.1670)$ and $(\$82,220, -0.0338)$, and continue iterating until we have sufficiently converged to the root of the signed distance. In this case, the root is $W(0) = \$79,900$.² In Figure 6, we have included the $W(0) = \$79,900$ EGPF, which, of course, goes through the desired point \mathbf{d} .

As a byproduct of the proximity algorithm, we also know that to obtain this desired point, we run the algorithm in

²The computational time to run this entire iterative process to obtain this root was 62 seconds on a Mac Studio computer with an M1 Max chip.

Appendix A with a utility of 100 assigned to Goal 1 and a utility of 108.47 assigned to Goal 2.³ This in turn generates the optimal investment strategy and the optimal goals-taking strategy that enables the initial investment of \$79,000 to obtain a 70% chance of attaining Goal 1 and an 80% chance of attaining Goal 2.

This subsection's algorithm for determining the minimum $W(0)$ to attain the investor's desired probabilities is particularly useful to investors who are well off. It enables them to devote the correct amount of funding to their goals, neither underspending so as to make it impossible to attain their desired probabilities nor, just as importantly, overspending, so we are able to optimally free their other funds for other financial priorities. For investors who are less well off, the proximity algorithm presented in the previous subsection will be of more use, enabling investors to use their money to come as close to their desired probabilities as is possible.

4.3. Incorporating Partial Goals

The algorithms in the previous two subsections can be extended to incorporate the important case of partial goals. As explained in the Introduction, a partial goal is an alternative to the full goal that is acceptable to the investor. While the partial goal costs less than the full goal, it also correspond to a lower utility when it is fulfilled in place of the full goal.

Consider the example in Table 1. Goal 1 is strictly a full goal that, 10 years from now, is either fulfilled for \$7,000 or not fulfilled. The investor wants at least an 80% chance of fulfilling this goal. Goal 2 is also a goal that the investor will either fulfill or forgo 10 years from now. The investor wants at least a 90% chance of fulfilling the full goal, which costs \$20,000, but they are also open to a partial goal that costs only \$9,000. They also want at least a 91% chance of fulfilling either the full or partial version of this goal. Goal 3, which the investor can be fulfill or forgo 15 years from now has three partial goal versions. Perhaps Goal 3 is a charitable donation goal, and the investor would like to give \$40,000 to the charity, but is also open to giving \$30,000, \$20,000, or \$10,000. The investor wants at least an 86% chance of giving one of these four amounts, at least an 85% chance of giving \$20,000 or more, at least an 80% chance of giving \$30,000 or more, and at least a 60% chance of giving the full \$40,000. Finally, Goal 4 is a long term goal in 30 years. The investor wants at least an 80% chance of fulfilling the full goal, which costs \$90,000 but is open to instead fulfilling the partial version of the goal, which costs \$50,000. The investor wants at least a 95% chance of fulfilling either the full or partial version of Goal 4.

If the investor only has \$20,000, they are unable to obtain

³Alternatively, since only the ratio of these utilities matter, any two utility values where 47.97% of the total utility is assigned to Goal 1 and the remaining 52.03% is assigned to Goal 2 can be used.

Goal (p = partial)	Time (in years)	Cost (in 1000s)	Desired minimum probabilities	Probabilities from p_{prox} if $W(0) = \$20K$	Probabilities from p_{prox} if $W(0) = \$39.4K$	Utilities for p_{prox} if $W(0)$ $= \$39.4K$
1	10	\$7	80%	56%	80%	100.00
2	10	\$20	90%	55%	91%	325.77
p	"	\$9	91%	82%	91%	130.82
3	15	\$40	60%	19%	60%	453.26
p	"	\$30	80%	44%	81%	368.92
p	"	\$20	85%	60%	87%	226.91
p	"	\$10	86%	74%	87%	82.94
4	30	\$90	80%	40%	85%	804.95
p	"	\$50	95%	70%	96%	431.22

Table 1. Results from determining the proximity point and the minimal initial wealth to obtain the investor's desired minimum probabilities for a case with partial, as well as full, goals.

their desired probabilities. The proximity point's probabilities are given in the third to last column in Table 1. We note that with partial goals, unlike full goals, the proximity point may satisfy some of the desired probabilities, but not all of them, although in this case it satisfies none of them. The algorithm from the previous subsection shows that to obtain all of the desired probabilities, the investor minimally needs \$39,400. The proximity point's probabilities when $W(0) = \$39,400$ are given in the second to last column in Table 1. Note for Goal 4, that the probability for fulfilling the full goal is 85%, which is five percentage points higher than the minimally required 80%. This overshoot of the minimally required probability is a byproduct of needing a 95% chance of fulfilling either the full or partial version of Goal 4, which is barely accomplished. Utility values that correspond to the proximity point when $W(0) = \$39,400$ are given in the last column of Table 1. Note that, unlike the probabilities, the utilities for partial goals are individual, not cumulative, so, for example, in Goal 3, the utility of 368.92 corresponds to the investor spending exactly \$30,000, not more or less. We observe that even though we only require the investor to specify desired goal probabilities (and not utilities for each goal), the algorithm is able to back out the investor's implied utility over goals, generating this latent information in a computationally facile manner.⁴

4.4. Sensitivity To Different Investment Portfolios

As stated in Subsection 3.1, our above examples have the investor optimally jump between 15 investment portfolios. These portfolios are on the Markowitz efficient frontier (EF) in Figure 1 for 15 equally spaced μ values from 5.26%, which corresponds to the vertex of the EF hyperbola, to 8.86%, the highest of the expected returns for the three index funds given in equation (4). The 3×3 return covariance for these three index funds is given in equation (5).

There will always be some error in forecasting the expected returns and covariances of these funds, so we explore the sensitivity of our results in the previous subsection to per-

⁴The computational time to determine all of the results in this table was 126 seconds on a Mac Studio computer with an M1 Max chip.

Efficient Goals Probabilities: A New Frontier

Goal (p = partial)	Time (in years)	Cost (in 1000s)	Desired minimum probabilities	Probabilities from p_{max} if $W(0) = \$20K$	Probabilities from p_{max} if $W(0) = \$33.6K$	Utilities for p_{max} if $W(0)$ $= \$33.6K$
1	10	\$7	80%	63%	82%	100.00
2	10	\$20	90%	65%	92%	320.68
p	"	\$9	91%	85%	92%	129.29
3	15	\$40	60%	30%	62%	419.94
p	"	\$30	80%	55%	81%	339.99
p	"	\$20	85%	69%	85%	197.29
p	"	\$10	86%	79%	86%	66.48
4	30	\$90	80%	53%	84%	637.68
p	"	\$50	95%	78%	96%	361.23

Table 2. The update to the results from Table 1 when we add one percentage point to the expected returns; that is, we replace μ from equation (4) with μ_{new} from equation (12).

turbing the values in both equation (4) and equation (5). If, for example, we add one percentage point to the expected returns in equation (4), we get

$$\mu_{\text{new}} = \begin{bmatrix} 0.0593 \\ 0.0870 \\ 0.0986 \end{bmatrix}. \quad (12)$$

Table 2 shows the effect of using μ_{new} in place of the old μ given in equation (4) on Table 1 from the previous subsection. Note that the 15 investment portfolios now correspond to 15 equally spaced μ values from 6.26%, which corresponds to the vertex of the new EF hyperbola, to 9.86%, the highest of the new expected returns for the three index funds. The probability of attaining each goal when $W(0) = \$20,000$ increases in the range of 3 to 13 percentage points due to using μ_{new} . Further, the minimum initial investment needed to attain all the desired goal probabilities decreases from \$39,400 to \$33,600.

On the other hand if we restore μ back to its original values in equation (4) but now increase the covariances 50% by multiplying Σ in equation (5) by 1.5, we get the new return covariance matrix

$$\Sigma_{\text{new}} = \begin{bmatrix} 0.0026 & -0.0026 & -0.0032 \\ -0.0026 & 0.0594 & 0.0463 \\ -0.0032 & 0.0463 & 0.0587 \end{bmatrix}. \quad (13)$$

Table 3 shows the effect of using Σ_{new} in place of the old Σ given in equation (5). The 15 investment portfolios have the same μ values as in the main examples, since the μ value for the vertex of the new EF hyperbola is still 5.26%. The 15 corresponding sigma values, however, are all increased, of course. The probability of attaining each goal when $W(0) = \$20,000$ generally decreases due to Σ_{new} , and by as much as seven percentage points, although the probability of attaining the first goal increases by four percentage points. The increased volatility causes the minimum initial investment needed to attain all the desired goal probabilities to increase from \$39,400 to \$42,900.

Because we have used the Markowitz efficient frontier to generate our 15 investment portfolios, we have allowed the possibility of short positions in any of the three component

Goal (p = partial)	Time (in years)	Cost (in 1000s)	Desired minimum probabilities	Probabilities from p_{max} if $W(0) = \$20K$	Probabilities from p_{max} if $W(0) = \$42.9K$	Utilities for p_{max} if $W(0)$ $= \$42.9K$
1	10	\$7	80%	60%	82%	100.00
2	10	\$20	90%	53%	92%	319.30
p	"	\$9	91%	79%	92%	126.38
3	15	\$40	60%	12%	63%	462.06
p	"	\$30	80%	41%	82%	366.54
p	"	\$20	85%	59%	86%	227.97
p	"	\$10	86%	74%	86%	86.43
4	30	\$90	80%	35%	86%	901.96
p	"	\$50	95%	64%	96%	473.49

Table 3. The update to the results from Table 1 when we increase the covariances by 50%; that is, we replace Σ from equation (5) with Σ_{new} from equation (13).

Portfolio	μ	Portfolio Weights (shorting)			Portfolio Weights (no shorting)		
		Index 1	Index 2	Index 3	Index 1	Index 2	Index 3
0	5.26%	0.9098	0.0225	0.0677	0.9098	0.0225	0.0677
1	5.52%	0.8500	0.0033	0.1467	0.8500	0.0033	0.1467
2	5.78%	0.7903	-0.0160	0.2257	0.7855	0.0	0.2145
3	6.03%	0.7305	-0.0352	0.3047	0.7201	0.0	0.2799
4	6.29%	0.6707	-0.0545	0.3837	0.6546	0.0	0.3454
5	6.55%	0.6110	-0.0737	0.4628	0.5891	0.0	0.4109
6	6.81%	0.5512	-0.0930	0.5418	0.5237	0.0	0.4763
7	7.06%	0.4915	-0.1122	0.6208	0.4582	0.0	0.5418
8	7.32%	0.4317	-0.1315	0.6998	0.3928	0.0	0.6072
9	7.58%	0.3719	-0.1507	0.7788	0.3273	0.0	0.6727
10	7.83%	0.3122	-0.1700	0.8578	0.2618	0.0	0.7382
11	8.09%	0.2524	-0.1892	0.9368	0.1964	0.0	0.8036
12	8.35%	0.1927	-0.2085	1.0158	0.1309	0.0	0.8691
13	8.61%	0.1329	-0.2277	1.0948	0.0655	0.0	0.9345
14	8.86%	0.0731	-0.2470	1.1738	0.0	0.0	1.0000

Table 4. The optimal weights for the three component indexes in the 15 investment portfolios if we allow short selling and then if we do not allow short selling.

index funds. Indeed, for 13 of these portfolios, the second index fund is optimally shorted, as can be seen in Table 4. If we wish to restrict ourselves to no short positions, then for each of the 15 portfolios' μ values, we can determine the optimal index fund weights using a constrained numerical optimizer. The resulting optimal nonnegative portfolio weights are contained in the last three columns of Table 4. Even though the weights seem considerably different if we prohibit short selling, the effect on the resulting 15 portfolio volatilities is rather small, as can be seen by comparing the Markowitz efficient frontier with the modified optimal frontier for no short selling seen in Figure 7. This means that prohibiting short selling has an almost negligible detrimental effect on Table 1. This is quantified by comparing Table 1 to Table 5, the results when short sales are prohibited. As an example, the minimum initial investment needed to attain all the desired goal probabilities increases when short sales are prohibited, as it must, but only from \$39,400 to \$39,700.

4.5. Infusions

Our dynamic programming algorithm in Appendix A accommodates any specified external infusions, $I(t)$, over time. For example, we again consider the example from Subsection 4.3 that leads to Table 1, but now, in addition to the initial wealth of $W(0) = \$20,000$, we will add an infusion stream from, say, a pension. Assume this pension

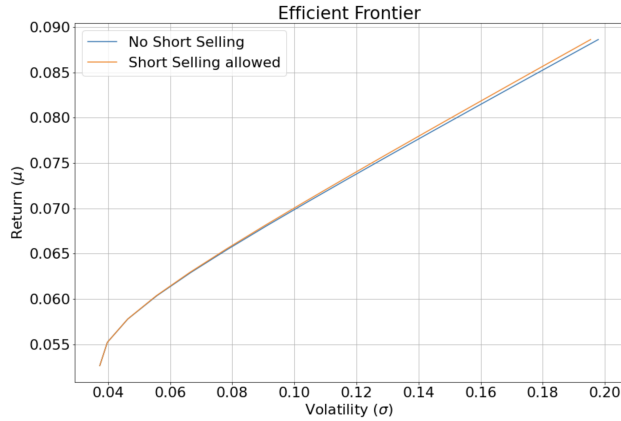


Figure 7. The Efficient Frontier (in orange), which allows shorting each of the three index fund positions, versus the modified frontier (in blue) if we are restricted from shorting any of the three positions. Note the small difference between the two frontiers.

Goal (p = partial)	Time (in years)	Cost (in 1000s)	Desired minimum probabilities	Probabilities from p_{prox} if $W(0) = \$20K$	Probabilities from p_{prox} if $W(0) = \$39.7K$	Utilities for p_{prox} if $W(0)$ $= \$39.7K$
1	10	\$7	80%	58%	80%	100.00
2	10	\$20	90%	55%	91%	325.50
p	"	\$9	91%	82%	91%	129.86
3	15	\$40	60%	19%	60%	452.10
p	"	\$30	80%	44%	82%	369.25
p	"	\$20	85%	60%	87%	228.69
p	"	\$10	86%	74%	87%	84.21
4	30	\$90	80%	40%	86%	826.94
p	"	\$50	95%	70%	97%	439.00

Table 5. The update to the results from Table 1 when we prohibit short selling. Note that the difference from Table 1 is quite small.

Goal (p = partial)	Time (in years)	Cost (in 1000s)	Desired minimum probabilities	Probabilities from p_{prox} if $W(0) = \$20K$	Probabilities from p_{prox} if $W(0) = \$33.7K$	Utilities for p_{prox} if $W(0)$ $= \$33.7K$
1	10	\$7	80%	57%	83%	100.00
2	10	\$20	90%	73%	92%	304.15
p	"	\$9	91%	86%	92%	123.56
3	15	\$40	60%	25%	60%	420.77
p	"	\$30	80%	55%	81%	337.15
p	"	\$20	85%	75%	87%	187.96
p	"	\$10	86%	80%	87%	53.75
4	30	\$90	80%	50%	87%	709.78
p	"	\$50	95%	83%	98%	370.92

Table 6. The update to the results from Table 1 when we add the effect of the infusions, $I(t)$, from the pension specified in equation (14).

starts in year 20 by paying \$2000, and then increases each year by a projected inflation rate of 3%; that is,

$$I(t) = \begin{cases} 0 & \text{if } t < 20 \\ 2 \times 1.03^{(t-20)} & \text{if } t \geq 20. \end{cases} \quad (14)$$

These infusions clearly will increase the optimal probability of attaining the goals when $t \geq 20$, but they also increase the probability of attaining the goals when $t < 20$, since knowledge of the later infusions frees the use of the initial investments for earlier goals. The beneficial effect of the infusions in equation (14) are shown in Table 6. Note that these infusions decrease the minimal initial investment needed to attain all the desired goal probabilities from \$39,400 to \$33,700.

Investors, of course, may not be able to find the additional \$39,400 (or \$33,700) at $t = 0$ needed to attain all their desired goal probabilities. Instead, they may wonder what minimal constant nominal annual contributions, $I(t) = c$ (such as payouts from a constant annuity), would allow them to attain all their desired goal probabilities in place of the additional upfront contribution. Or they may want the minimal constant real annual contributions,

$$I(t) = c \times 1.03^t, \quad (15)$$

assuming, as before, a 3% inflation rate.

The method in Subsection 4.1 can easily be adapted to answer this question by replacing the role of $W(0)$ in that subsection with the role of c in these infusion models. Applying this adapted method to the example from Subsection 4.3, which leads to Table 1 where $W(0) = \$20,000$, we find that $c = 1.05$ is the minimal value in equation (15) needed to attain all the desired goal probabilities. The corresponding version of Table 1 is shown in Table 7.

For many investors, these payments over time represent a more realistic option than investing more money initially. These single factor models for paying over time, however, can lead to infusions that are unnecessarily high, as is the case here. The infusion model in equation (15) is not optimal for our problem in that c must be high enough to ensure attaining the early goals' desired probabilities, but because

Goal (p = partial)	Time (in years)	Cost (in 1000s)	Desired minimum probabilities	Probabilities from p_{prox} if $c = 0$	Probabilities from p_{prox} if $c = 1.05$	Utilities for p_{prox} if $c = 1.05$
1	10	\$7	80%	56%	81%	100.00
2	10	\$20	90%	55%	91%	301.51
p	"	\$9	91%	82%	91%	96.81
3	15	\$40	60%	19%	63%	429.23
p	"	\$30	80%	44%	82%	337.41
p	"	\$20	85%	60%	92%	192.04
p	"	\$10	86%	74%	92%	52.72
4	30	\$90	80%	40%	89%	799.40
p	"	\$50	95%	70%	99%	413.33

Table 7. The update to the results from Table 1 when we add the effect of infusions over time of the form $I(t) = c \times 1.03^t$. The initial investment, $W(0) = \$20,000$.

it maintains the same infusion in real dollars in later years, it overfunds the final goals. This can be observed by noting the unnecessarily high probabilities attained for the $t = 30$ goals in Table 7, which implies that the infusions after $t = 15$ can be reduced while still maintaining the desired goal probabilities.

5. Concluding Comments

This paper introduces a new efficient frontier in the context of goals-based wealth management called the efficient goals probability frontier (EGPF). Just as the Markowitz (1952) mean-variance efficient frontier presents the optimized trade-off between risk and return in investment portfolios, the EGPF presents the optimized trade-offs between the probabilities of attaining two or more competing goals that occur at various times with varying costs. While the Markowitz (1952) mean-variance efficient frontier resides in a static framework, the EGPF lives in a dynamic optimization framework.

The EGPF enables investors to select the specific set of trade-offs — which corresponds to a specific point on the EGPF — that best fits their preferences. Behind each point on the Markowitz efficient frontier is a static portfolio, whereas behind each point on the EGPF is a set (i.e., a vector) of utility values for each of the goals, which corresponds to the optimal dynamic investment strategy for trading between investment portfolios over time and the optimal goals-taking strategy for determining which goals an investor is best off fulfilling or forgoing. These utility values are intentionally behind the scenes in the EPGF, since investors intuitively understand the concept of the probabilities of attaining their goals, not the utility values assigned to their goals.

In contrast to the Markowitz frontier, the EGPF has limitations, despite the advantages mentioned above: (i) Unlike the mean-variance efficient frontier, the EGPF is not available in closed form, because generating it requires numerical optimization. (ii) Since, for n goals, the EGPF is an $n - 1$ dimensional hypersurface in n dimensional probability space, it is hard to visualize graphically when the number of goals, n , exceeds three. (iii) Generating the EGPF is computation-

ally heavy, as the multiple goals probability optimization has to be repeated for all chosen utility weight vectors, which can take a few minutes of computational time.

The intuitions provided by the EGPF, however, can be exploited to determine three key results that are not affected by the limitations in (ii) and (iii) above. Each of these three key results requires the investor to specify desired probabilities for attaining each of their goals. The first key result is the ability to determine a set of attainable probabilities that are close to the desired probabilities. These attainable probabilities correspond to a specific point on the EPGF for the initial investment worth. The second key result is the ability to determine the minimum initial investment necessary to obtain each of the desired probabilities. The third key result is the ability to determine the minimum of a one-parameter family of infusions necessary to obtain each of the desired probabilities. For example, we can determine the minimum yearly constant infusions in either nominal dollars or in real dollars necessary to obtain the desired probabilities. In each of these three cases, we know the utilities associated with these key results and can therefore determine the optimal investment strategy and the optimal goals-taking strategy to obtain them. Further, these three key results can be determined when we have partial goals, as well as full goals, such as when an investor is okay with buying cheaper car models if that is wiser than purchasing the fancy car model they might ideally want.

In future work, we look to explore the question of delaying the purchase of a goal to a better year. At the moment, each goal has a time associated with it, and the investor must purchase the goal (or a partial version of the goal) at that time or decide to forgo it completely. But investors may wish instead to simply put off the goal to a later year. Because the underlying analysis in this paper uses dynamic programming, which evolves information backwards in time, it is a poor tool to use with delayed goal-taking, which is a forwards-in-time concept. Recent approaches connected to machine learning, i.e., deep reinforcement learning, show hope for making progress with this difficult but important problem.

References

- Alexander, G. J., A. M. Baptista, and S. Yan (2017, March). Portfolio selection with mental accounts and estimation risk. *Journal of Empirical Finance* 41, 161–186.
- Bellman, R. (1952, August). On the Theory of Dynamic Programming. *Proceedings of the National Academy of Sciences* 38(8), 716–719.
- Brown, D. C., S. Cederburg, and M. S. O'Doherty (2017, December). Tax uncertainty and retirement savings diversification. *Journal of Financial Economics* 126(3), 689–712.

- 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.
- Consiglio, A., F. Cocco, and S. A. Zenios (2004). www.Personal.asset.allocation. *Interfaces* 34(4), 287–302. Publisher: INFORMS.
- Das, S., H. Markowitz, J. Scheid, and M. Statman (2011, July). Portfolios for Investors Who Want to Reach Their Goals While Staying on the MeanVariance Efficient Frontier. *The Journal of Wealth Management* 14(2), 25–31.
- Das, S. R., D. Ostrov, A. Casanova, A. Radhakrishnan, and D. Srivastav (2021, July). Combining Investment and Tax Strategies for Optimizing Lifetime Solvency under Uncertain Returns and Mortality. *Journal of Risk and Financial Management* 14(7), 285. Number: 7 Publisher: Multidisciplinary Digital Publishing Institute.
- 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 (2020). Dynamic Portfolio Allocation in Goals-Based Wealth Management. *Computational Management Science* 17(June), 613–640.
- Das, S. R., D. Ostrov, A. Radhakrishnan, and D. Srivastav (2021, May). Dynamic optimization for multi-goals wealth management. *Journal of Banking & Finance*, 106192.
- Das, S. R. and G. Ross (2020, September). The Role of Options in Goals-Based Wealth Management. SSRN Scholarly Paper ID 3691699, Social Science Research Network, Rochester, NY.
- Das, S. R. and M. Statman (2013). Options and structured products in behavioral portfolios. *Journal of Economic Dynamics and Control* 37(1), 137–153.
- 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.
- 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., 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 (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.
- Martellini, L., V. Milhau, and J. Mulvey (2019, June). Flexicure Retirement Solutions: A Part of the Answer to the Pension Crisis? *The Journal of Portfolio Management* 45(5), 136–151. Publisher: Institutional Investor Journals Umbrella.
- Martellini, L., V. Milhau, and J. Mulvey (2020, April). Securing Replacement Income with Goal-Based Retirement Investing Strategies. *The Journal of Retirement* 7(4), 8–26. Publisher: Institutional Investor Journals Umbrella.
- 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.
- Merton, R. C. and A. Muralidhar (2017, April). Time for Retirement SeLFIES? SSRN Scholarly Paper ID 2945668, Social Science Research Network, Rochester, NY.
- Muralidhar, A. (2018, December). The F-Utility of Wealth: Its 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.
- Parker, F. J. (2020, April). Allocation of Wealth Both Within and Across Goals: A Practitioners 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. 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.
- Vo, H. T. and R. Maurer (2013, March). Dynamic Asset Allocation with Regime Shifts and Long Horizon CVaR-Constraints. SSRN Scholarly Paper ID 2191286, Social Science Research Network, Rochester, NY.
- 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.

Appendices

A. Algorithm For Dynamic Optimization With Multiple Goals

Large-scale dynamic optimization with hundreds of goals over decades is elaborated in [Das, Ostrov, Radhakrishnan, and Srivastav \(2021\)](#). That paper’s model considers each of n goals, where, for each goal, we know the time at which the goal can be fulfilled, the cost of fulfilling the goal, and the utility that fulfilling a goal provides to the investor. We also assume that we know all infusions, $I(t)$, such as money received from Social Security or an annuity, over time. The intent is to maximize the sum of the utility values from the fulfilled goals.

This is done by optimizing over both a finite choice of goals-taking strategies, indexed by k , and a finite choice of available investment portfolios, indexed by l . A goals-taking strategy states which goals at a given time the investor will choose to fulfill versus forgo given their portfolio worth, W , at that time. Each goals-taking strategy at a time t has an associated cost, $c_k(t)$, representing the sum of the

costs of the goals to be fulfilled according to the strategy. The available investment portfolios have known expected returns, μ_l , and volatilities, σ_l , which can correspond to locations on the efficient frontier, such as in Figure 3, or locations below the frontier.

This model also accommodates partial goals. That is, the investor may ideally have a goal of purchasing an expensive car, but also be open to partial goals, meaning less expensive cars that also bring less utility to the investor.

We briefly summarize the approach here.

A.1. Basic Setup

We solve the multiple goals optimization problem using dynamic programming on a grid of portfolio wealth and time, i.e., (W, t) . For this paper’s examples, the values of time, $t = 0, 1, 2, \dots, T$, are in years. The values of W lie in a range of possible wealth values determined by geometric Brownian motion, which, for a given goals-taking strategy, k , and a given portfolio investment, l , is given by:

$$W(t+h) = [W(t) + I(t) - c_k(t)]e^{\left(\mu_l - \frac{\sigma_l^2}{2}\right)h + \sigma_l\sqrt{h}Z}, \quad (16)$$

where $h = 1$ (year) in the examples in this paper.

We determine $[W_{min}, W_{max}]$, the possible range for $W(t)$, by determining the lowest and highest possible values of W over time through applying equation (16) with, as appropriate, the smallest or largest possible values for c_k , μ_l , σ_l , and Z , assuming that, realistically, $|Z| \leq 3$. We then place the interior grid points for W between W_{min} and W_{max} so that the natural logarithm of their values are equally spaced apart. This spacing offers greater accuracy and stability for the solution.

A.2. Transition Probabilities

Define $\phi(z)$ to be the value of the probability density function of the standard normal random variable at $Z = z$. Assume we start at time period t with $W_i(t) + I(t) - c_k(t)$ dollars and progress a year to the time period $t + 1$. Here, $W_i(t)$ is the wealth amount at node i on the solution grid at time t . We are interested in computing the probability of transitioning from $W_i(t)$ to node $W_j(t+1)$, i.e., the wealth level at node j in the next year. Rearranging equation (16) with $h = 1$, we obtain the relative transition probabilities, \tilde{q} , from $W_i(t)$ to $W_j(t+1)$:

$$\tilde{q}(W_j(t+1)|W_i(t)) = \phi\left(\frac{\ln\left(\frac{W_j(t+1)}{W_i(t) + I(t) - c_k(t)}\right) - \left(\mu_l - \frac{\sigma_l^2}{2}\right)}{\sigma_l}\right). \quad (17)$$

These are relative because their discrete sum over all the j

nodes at time $t + 1$ is not necessarily equal to one. Therefore, we normalize the relative transition probabilities, \tilde{q} , to obtain q , our transition probabilities, which sum to one:

$$q(W_j(t+1)|W_i(t)) = \frac{\tilde{q}(W_j(t+1)|W_i(t))}{\sum_j \tilde{q}(W_j(t+1)|W_i(t))}. \quad (18)$$

The transition probabilities, q , are used for dynamic programming and provide the probability of going from the wealth level $W_i(t)$ at time t to each possible wealth level $W_j(t+1)$ at time $t+1$. Note that, from equation (17), the transition probabilities depend on the chosen investment portfolio, l , at $W_i(t)$.

A.3. Backwards Recursion

We are now ready to solve the dynamic optimization problem in which we look to maximize the expected utility aggregated over all present and future goals, which we denote as the value function. More specifically, the value function, $V(W_i(t), t)$, at time t at a given level of wealth, W_i , on the grid is the highest possible expected sum of the utilities from fulfilled goals from time periods $\geq t$, given that the investor starts with wealth W_i at time t . In obtaining the value function at a grid point, (W_i, t) , we must also determine the optimal goals-taking strategy, k^* , and the optimal investment strategy, l^* , at this grid point.

To determine V , k^* , and l^* at each grid point $(W_i(t), t)$, we apply the Bellman (1952) equation,

$$V(W_i(t), t) = \max_{k, l} \left[U_k(t) + \sum_j V(W_j(t+1)) \cdot q(W_j(t+1)|W_i(t)) \right], \quad (19)$$

to solve the dynamic programming problem. The value of $U_k(t)$ in equation (19) is the combined utility of the goals that are chosen to be fulfilled at time t using goals-taking strategy k . The values of k^* and l^* are the maximizing k and l values on the right-hand side of the Bellman equation (19). The Bellman equation is applied backwards in time: At time T , for each $W_i(T)$, we can easily determine V and k^* (there is no l^* at time T) by selecting the goals-taking strategy k where U_k is highest, subject to $c_k \leq W_i$. We can then use the Bellman equation to determine V , k^* and l^* at each grid point $(W_i(T-1), T-1)$ (that is, each wealth level W_i when $t = T-1$), then each grid point $(W_i(T-2), T-2)$, etc., until we finish with each grid point $(W_i(0), 0)$. At the end, we have $V(i, t)$, $k^*(i, t)$, and $l^*(i, t)$; that is, we have the value function, the optimal goals-taking strategy, and the optimal investment portfolio at every $(W_i(t), t)$ grid point.

A.4. Forward Calculation Of Goal Probabilities

Once we have the $q(W_j(t+1)|W_i(t))$, $k^*(i, t)$, and $l^*(i, t)$, we can use them to determine the unconditional probability

distribution, $p(W_i(t), t)$, for the investor's wealth at each time t . To do this, we evolve the unconditional probability distribution forward in time, starting with $t = 0$, then $t = 1$, and ending with $t = T$. This is called the "forward pass" that delivers the unconditional probability of reaching any level of wealth at any time, given that we know the investor's initial wealth, W_0 , at time $t = 0$.

More specifically, at $t = 0$, $p(W_0(0), 0) = 1$ and $p(W_i(0), 0) = 0$ for all $W_i \neq W_0$. Given that we know $p(W_i(t), t)$, the unconditional wealth probability distribution at time t for each i , we can obtain the distribution at time $t+1$ by using the following "forwards equation," which we can apply to each wealth level W_j :

$$p(W_j(t+1), t+1) = \sum_i q(W_j(t+1)|W_i(t)) \cdot p(W_i(t), t) \quad (20)$$

for $t = 0, 1, 2, \dots, T-1$.

Once this recursive application of equation (20) ends, we have $p(W_i(t), t)$, the unconditional probability distribution for every point in the wealth grid at every time period of the portfolio.

Combining this $p(W_i(t), t)$ with $k^*(i, t)$, we can determine the probability of fulfilling each goal or partial goal at any given time t by summing $p(W_i(t))$ over each W_i node where the goal is taken according to the optimal goals-taking strategy $k^*(i, t)$.

A.5. Relationship To The Efficient Goal Probability Frontier

The system diagram shown in Figure 8 provides a high-level depiction of the process described here in the backward recursion and forward pass steps in the algorithm. Overall, the algorithm maps a set of goals whose times, costs, and utilities are known to a set of corresponding optimized goal probabilities. In reality, the times and the costs of the goals would be known, but not the utilities. One set of utilities leads to a single point on the EGPF (efficient goal probability frontier). By varying the utilities for the goals, we can trace out the EGPF.

Because the EGPF only displays probabilities, not utilities, it works in a language investors can understand. Once they see the trade-offs between the probabilities of attaining their goals, they can select the best trade-off for their preferences, meaning the best point for them on the EGPF. This then specifies the set of utilities, and allows us to compute the optimal goals-taking strategy and the optimal investment strategy tailored to the investor's preferences using this Appendix's algorithm.

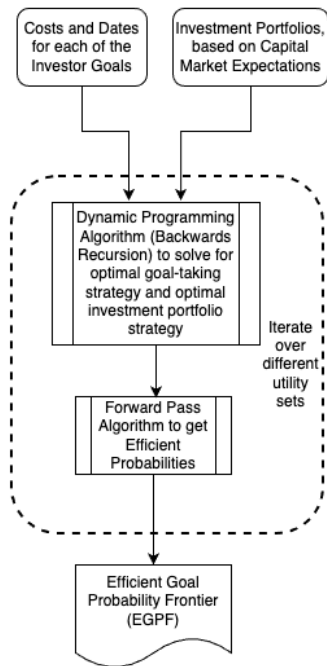


Figure 8. Process flow of the algorithm. The investor inputs the dates for when each of their goals can be fulfilled, along with the cost of fulfillment. The algorithm then uses possible investment portfolios based on capital market expectations to compute both the optimal goals-taking strategy and the optimal investment strategy by solving the Bellman equation for dynamic programming. These optimal choices are then injected into the forward pass of the algorithm to get one point on the EGPF. By rerunning this for a selection of utility sets for the goals, the full EGPF is generated.