Choosing and Using Utility Functions in Forming Portfolios

Geoffrey J. Warren

Geoffrey J. Warren is an associate professor at the College of Business and Economics, Australian National University, Canberra, Australia.

Utility functions offer a means to encode objectives and preferences in investor portfolios. The functions allow one to place a score on outcomes and then identify optimal portfolios by maximizing utility. The central theme of this article is that utility functions should be tailored to the investor. I discuss how an appropriate function might be chosen and demonstrate concepts for power utility and reference-dependent utility. A modeling approach is presented that may be applied without resorting to dynamic optimization. The selection of utility functions is illustrated for four investor types.

Editor's Note: The original version had a typographical error in Figure 4B, which has been corrected in this version.

Disclosure: The author reports no conflict of interest.

CE Credits: 1

Utility functions are common in the academic literature but not widely used for forming portfolios in practical settings. Mean-variance analysis remains the dominant approach employed by the investment industry, in part because it is familiar. Although mean-variance analysis generates useful results in certain situations, a number of features limit its suitability in other situations. First, mean-variance analysis is designed to analyze portfolio return and risk over a discrete (and often limited) horizon. *Variance* is a measure of the variability of investment outcomes over a specific interval; monthly, quarterly, and yearly intervals are typical, although sometimes the horizon is extended to span multiple years. This limitation can be a problem for long-term investors, especially those who intend to draw on wealth over many periods and hence do not operate within one horizon. Second, mean-variance analysis may not align with investor objectives, especially those involving a target, such as delivering a real return, a required income stream, or sufficient assets to cover liabilities. Finally, mean-variance analysis may fall short in cases where outcomes are highly skewed, such as the evaluation of wealth over long horizons, or where option-like payoffs are involved. Consider an individual investing to generate an adequate income in retirement. Analyzing the mean and variance of his or her portfolio returns over some discrete time horizon is only vaguely relevant to the main concern—namely, the stream of income that can be drawn over time.

This article explores the use of utility-based methods for forming portfolios. A central theme is that the utility function and related analysis should be tailored to reflect the investor's circumstances. A wide range of investment situations can be effectively analyzed by modeling the evolution of total portfolio value—or accumulated wealth—over the horizon of concern and using a purposefully selected utility function to score the associated outcomes. These outcomes may be withdrawals made from the portfolio and/or the portfolio value itself. By structuring the analysis to generate a distribution of outcomes that depends on portfolio weights, we can compare different portfolios or identify optimal portfolios by finding the asset mix that maximizes expected utility across all potential outcomes. I outline here an approach for using utility functions to identify optimal portfolios that can be readily applied

in practice without delving into overly complex mathematical or simulation techniques. I illustrate the selection of utility functions and application of the approach for four investor types.

Proposing the use of utility functions for forming portfolios is not new. Adler and Kritzman (2007) made a similar appeal. A number of authors have used utility functions to evaluate portfolios in an applied context. For instance, retirement outcomes were analyzed by Blake, Wright, and Zhang (2013), who used a variation of the value function from the prospect theory of Tversky and Kahneman (1992); by Levy (2016), who used power, log, and prospect theory utility; by Bell, Liu, and Shao (2017a, 2017b), who used power utility plus a bequest term; and by Estrada and Kritzman (2018), who proposed a kinked utility function reflecting power utility above target consumption and a linear penalty below. The present study contributes by demonstrating how a utility-based approach might be implemented by industry practitioners to cater to a wide variety of investors with specific objectives and preferences. “The investor” might be an individual, a cohort of individuals, or an organization.

Choosing a suitable utility function is essentially a subjective task. It may be approached by (1) forming an understanding of the investor’s objectives and preferences over *outcomes* and then (2) selecting a functional form and parameters to represent the investor’s situation. The process can be assisted by gauging the proposed function’s implications. Caution is required when deviating from the academic literature, which is typically careful in its proposal and testing of utility functions. Nevertheless, the literature has reached no consensus about which utility function provides the “best” description of individual behavior (see Starmer 2000). Papers also tend to focus on individuals rather than organizations. A key recommendation here is to turn the process on its head by starting with the investor and then choosing a utility function to suit. Utility functions are not one-size-fits-all, which implies that the functional form is somewhat up for grabs. Furthermore, I demonstrate that perverse implications may arise from extracting a functional form and parameters from the academic literature without checking that the function is appropriate for the investor. Choosing a utility function that is fit for a purpose seems more important than seeking validation from an unsettled literature.

I take no stand on which is the “right” utility function in any particular situation. The aim here is to present

the case that a utility function should be purposefully selected, while recognizing that investors differ. I discuss the process by which a utility function might be chosen and underline the importance of the choice. Power utility and reference-dependent utility are used for demonstration. Reference-dependent utility is presented in two functional forms—one defined by the *difference* between the projected and target outcomes and the other defined by the *ratio* of these outcomes. These utility functions are broadly used within the related finance literature. Together, they provide considerable flexibility to accommodate a wide range of investors.

I also discuss the implications of utility function choice for portfolio formation, in particular how the preference for equities is affected by horizon and distance from the target. I show how power utility tends to constrain the preference for equities as the horizon lengthens because a heavy penalty is imposed on the extreme wealth losses that may occur over long periods. In contrast, reference-dependent utility tends to favor equities as horizon lengthens, which reflects both a tendency to penalize the lower tail relatively less aggressively and the fact that the higher returns on equities shift the distribution upward over time. Another implication of reference-dependent objective functions is that the optimal equity weighting often emerges as a U-shaped function of distance from the target.

The outline of a basic approach for forming portfolios is used to demonstrate how utility-based analysis can be readily applied and can generate plausible results. The approach entails four steps:

1. *Choose a utility function.* Select a “tailored” utility function and parameterize it to represent the objectives and preferences of the investor.
2. *Specify a plan for withdrawal and evaluation.* Specify a predetermined plan for any withdrawals from the portfolio as well as the combination of withdrawals and terminal portfolio value to be evaluated. Specifying the plan also entails setting the investment and, hence, analysis horizon.
3. *Project joint distributions of wealth outcomes for candidate assets.* Select the assets to be considered, and then project a joint distribution of expected wealth from investing in each asset over the horizon. This step is done by generating accumulated wealth outcomes for each asset associated with a range of paths (or, alternatively, scenarios or states).

4. *Solve for optimal portfolio weights.* Optimal portfolios are identified as follows:
 - A series of withdrawals and portfolio values is generated for each path by projecting forward. The closing portfolio value in each period is estimated as a function of the opening portfolio value, the asset weights, the change in accumulated wealth associated with each asset from step 3, and any withdrawals specified under step 2.
 - Within each simulated path, the withdrawals and/or terminal portfolio value are scored by using the utility function defined in step 1, in accordance with the evaluation plan from step 2.
 - Utility estimates are averaged (probability-weighted) across all paths to generate expected utility.
 - The optimal asset weights are identified as those that maximize expected utility.

These steps generate optimal asset weights over the horizon for the particular utility function, a predetermined plan for withdrawal and evaluation, and the joint asset distributions—as specified in steps 1, 2, and 3. The approach is relatively straightforward to apply and may be handled within a spreadsheet package such as Microsoft Excel. Complexity is reduced by restricting the scope of analysis, in particular by *predetermining* the withdrawal strategy and solving for *static* asset weights. Reducing complexity in this way means that the approach does not immediately deliver an optimal *dynamic* solution for the investment and withdrawal strategy. The issue of when dynamic analysis of some form might be appropriate is discussed later in the article.

In the illustrative analysis presented here, optimal portfolios are identified for four representative investor types—a private investor, an endowment fund, a defined-benefit fund, and a retired individual. I also provide an extension in which candidate investment strategies are compared on the basis of the expected utility they generate. The analysis demonstrates the flexibility of utility-based approaches while highlighting the importance of the choice of utility function, including its parameterization and the horizon over which it is applied.

This study proceeds by initially discussing the use of utility functions and then addressing their selection and parameterization. The approach for forming portfolios is subsequently outlined and illustrated

for the four representative investors. The use of utility measures to compare candidate strategies is then demonstrated. The question of static versus dynamic analysis is considered before reaching conclusions.

Rationale for Using Utility Functions

The advantages of using utility functions to form portfolios largely relate to them providing a mechanism for directly evaluating *all potential outcomes*. They do this by attaching a score to each point in the distribution. Expected utility can then be estimated as the probability-weighted average of these scores. Portfolios can be evaluated by the overall expected utility scores they deliver.

The main alternative to using utility functions involves evaluating portfolios based on *metrics* that summarize the distribution of outcomes. Mean–variance analysis is one such approach, where portfolios are often selected based on Sharpe ratios. Portfolios may be evaluated by using other metrics, such as probability of shortfall, value at risk, conditional value at risk, drawdown, and failure rates.

A key advantage of utility-based approaches is that they capture all information about the range of potential outcomes, whereas metrics may discard some relevant information. Adler and Kritzman (2007) argued that a utility-based approach (they called it “full-scale optimization”) is more comprehensive and relevant than defining investor preferences in relation to mean and variance. Consistent with this concept, Goetzmann, Ingersoll, Spiegel, and Welch (2007) highlighted the shortcomings of mean–variance metrics, such as Sharpe ratios and alpha, in the context of performance evaluation. They demonstrated that “manipulation-free” measures should be anchored in investor utility.

A related advantage is that utility functions accommodate the analysis of distributions of any shape. This attribute can be particularly relevant over long horizons, when the distribution of portfolio outcomes is often positively skewed because wealth is bounded by zero on the downside but unconstrained on the upside. This argument extends to analysis of option-like payoffs. Rubinstein (1976) argued that specifying the form of the utility function is superior to imposing assumptions about asset returns under many applications.

Utility-based methods offer considerable flexibility. A variety of functions are available that can encapsulate a wide range of investor objectives and preferences. They can be designed to recognize and place appropriate scores on various parts of a distribution. Unlike shortfall measures, utility functions can directly give appropriate credit for the upside. And if the investor has a particular dislike for being in the lower tail of the distribution, the utility function may be specified to ensure that lower-tail outcomes are sufficiently penalized. Utility functions can also be applied to multiperiod outcomes because utility scores can be added up over time. This characteristic can be particularly helpful where the objective relates to a stream of withdrawals from the portfolio, such as generating income in retirement. Indeed, the pension industry worldwide faces the task of developing better retirement products that combine an investment strategy and a withdrawal strategy, most likely in a dynamic framework. It is hard to imagine how the mean-variance techniques of modern portfolio theory can provide the required toolkit. Utility functions offer the analytical mechanics to move forward in this key area.

Utility-based approaches are correctly subject to the criticism that choosing a utility function is difficult and subjective. Analysis based on metrics, however, requires a mechanism to trade off the expected outcome against a risk measure. This trade-off involves addressing investor preferences and reflecting them in a risk-aversion parameter—or some rule of thumb that dictates how return is traded off against risk. An underlying utility function may be implied in the process even if it is not explicitly specified. Some degree of subjective choice is unavoidable either way. The ultimate question is which approach is more effective at addressing the problem facing the investor. In many circumstances, a utility-based approach will work better than relying on metrics.

Finally, utility-based approaches have broader uses than portfolio optimization. Utility functions can be used to place an overall score on a proposed portfolio, strategy, or action with the aim of establishing whether value is added relative to the status quo. For instance, the potential addition of a new asset class to a portfolio might be evaluated in terms of whether it increases utility at the margin. Utility functions can also be used to choose between alternatives. Utility-based analysis facilitates such decisions by summarizing value in a single number. The approach is similar to comparing portfolios or investment strategies by, for instance, their Sharpe ratios or information ratios.

Moreover, utility may be converted into economically meaningful measures, such as certainty equivalents, to aid comparisons.¹

Three Utility Functions

I take no stand on what is the “correct” utility function. Indeed, my key message is that the utility function should be selected and parameterized to reflect investor circumstances. This section analyzes three utility functions: power utility and two variations of reference-dependent utility. These functions are chosen because they are widely used, are capable of addressing many practical situations, and suffice for the purpose of demonstrating the main concepts and conveying how utility functions might be interpreted. A number of utility functions are available; the general family known as hyperbolic absolute risk-aversion (HARA) utility and Epstein–Zin–Weil utility are worth noting.²

Power utility, which lies within the HARA family, is broadly used within the academic literature. Equation 1 describes the functional form:

$$U_{PU,t} = \frac{W_t^{(1-CRRA)}}{1-CRRA} \quad (1)$$

where

$U_{PU,t}$ = power utility, period t

W_t = outcome, period t

CRRA = coefficient of relative risk aversion

[If CRRA = 1, $U_{PU,t} = \ln(W_t)$, i.e., log utility.]³

Power utility is defined over the *level* of the outcome being evaluated. Equation 1 specifies this outcome as W , which alludes to wealth, but consumption and income levels are other common choices. Power utility assumes constant relative risk aversion, which means that the curvature of the function is the same regardless of the level of the outcome.⁴ Power utility may thus be appropriate for an investor who is concerned with comparing outcomes on the basis of their level and spread but whose risk aversion does not change with the overall level.

A reference-dependent utility function evaluates outcomes relative to some reference level. Such functions may be more suitable than power utility for an investor with an objective that entails achieving some target. Various authors have proposed

functions that define utility with reference to a wealth or return target—for example, Hogan and Warren (1972), Fishburn (1977), Holtzhausen (1981), Stutzer (2003), Kahneman and Tversky (1979), Tversky and Kahneman (1992), Anthonisz (2012), and Estrada and Kritzman (2018). Furthermore, experimental studies support the idea that many investors perceive risk as related to shortfall versus some reference point (see Unser 2000 and Veld and Veld-Merkoulova 2008). A number of authors have used utility functions that reflect the value function component in the prospect theory of Kahneman and Tversky and Kahneman in evaluating investment outcomes⁵—for example, Bierman (1998), Adler and Kritzman (2007), Kőszegi and Rabin (2007), Blake et al. (2013), and Levy (2016). I follow this approach here, while noting that various other forms of reference-dependent utility exist, including habit persistence (see, e.g., Campbell and Cochrane 1999 and Grishchenko 2010) and the inclusion of floor levels (see Kingston and Thorp 2005).⁶

A reference-dependent utility function defined over the *difference* between projected outcome and target outcome is described by Equation 2, and one defined over the *ratio* appears as Equation 3.

The functional form of Equation 2 is similar to the value function in prospect theory, except that it is generalized here by providing for a weighting parameter, γ , to be applied to gains above the target. Under prospect theory, the parameter γ is set equal to 1. The reference-dependent utility function in ratio form, Equation 3, follows Tarlie (2017).

The reference point is the target outcome, W^* . Note that both equations provide for the target to vary with time. The deviation between projected and target outcome is moderated by curvature parameters α and β and multiplied by weighting parameters γ and λ . Different curvature and weighting parameters may be applied in the realms of outcomes above and below the target, providing scope to discount gains and penalize losses to varying degrees of intensity and in a nonlinear manner if desired. Both functions are positive when $W > W^*$, negative when $W < W^*$, and equal to zero when $W = W^*$.

An implication of incorporating a reference level is that the evaluation depends on where the overall distribution of outcomes lies with respect to the target. The reason is that the function is kinked; hence, the way that a given distribution is evaluated and scored is no longer independent of the outcome level. As

Function	Utility if $W > W^*$ (gain)	Utility if $W < W^*$ (loss)	Utility if $W = W^*$ (at target)	
$U_{RDU(D),t}$	$\gamma \left[(W_t - W_t^*)^\alpha \right]$	$-\lambda \left[(W_t^* - W_t)^\beta \right]$	0	(2)
$U_{RDU(R),t}$	$\gamma \left[\left(\frac{W_t}{W_t^*} \right)^\alpha - 1 \right]$	$\lambda \left[\left(\frac{W_t}{W_t^*} \right)^\beta - 1 \right]$	0	(3)

Where

$U_{RDU(D),t}$ = reference-dependent utility in difference form, period t

$U_{RDU(R),t}$ = reference-dependent utility in ratio form, period t

W_t = outcome, period t

W_t^* = target outcome, period t

α = curvature parameter on gains

β = curvature parameter on losses

γ = weighting parameter on gains

λ = weighting parameter on losses

a given distribution of outcomes is scaled up above the target, a diminishing portion of that distribution is evaluated as a loss and an increasing portion is evaluated as a gain. Therefore, less of the distribution is attracting the penalty applied to losses, which is steeper than the discount that is applied to gains. This feature turns out to be important as the horizon lengthens when assets offer different levels of return, as will be demonstrated.

Finally, reference-dependent utility functions of the general type presented in Equation 2 and Equation 3 have some practical advantages. One is that four parameters provide considerable flexibility. Tarlie (2017) discussed how his function embeds a number of utility functions discussed in the literature. In

addition, utility estimates arrive in intuitive units, being either dollar numbers or percentage deviations from the target, which can be helpful in calibrating the parameters. In contrast, power utility has just one parameter, and the utility numbers are largely uninterpretable.⁷ Nevertheless, one needs to be careful not to overplay this feature. Utility functions operate only as scores to rank outcomes (and portfolios). They need to be converted into certainty equivalents or something similar to become economically meaningful.

Functional Forms Compared: Power Utility vs. Reference-Dependent Utility

Utility. Charting utility functions to gauge how they evaluate outcomes can be helpful in choosing an appropriate function. **Figure 1** presents two charts that illustrate the functional forms of power utility (Panel A) for CRRA of 4 and reference-dependent utility (Panel B) in ratio form using the parameters of Blake et al. (2013). The charts span wealth outcomes ranging from 0.2 to 2.4, with target wealth set at 1.0 (allowing for the same x-axis).

The aim here is to characterize each function (the implications of parameter choice are investigated in the next section). Figure 1 highlights that the power utility function is a smooth curve but is *relatively* harsh on outcomes in the lower end of the range. Indeed, power utility provides comparatively little differentiation for outcomes that are near or above wealth of 1.0, where the slope is nearly indiscernible (at least for CRRA = 4). The reference-dependent utility function is kinked around the target. This kink reflects loss aversion; that is, the penalty applied to losses (shortfall) in relation to the target is much greater than the discount that is applied to gains of an equivalent magnitude. The curvature implies far more “evenly paced” changes in utility scores, however, than under power utility as outcomes move away from the reference point. Most notably, while scores under power utility (for CRRA > 1) decline more than exponentially as wealth decreases, the reference-dependent function declines in something much closer to a linear manner, with the weighting parameter of 4.5 dominating.

The key implication to draw from Figure 1 is that power utility has a propensity to penalize outcomes in the lower tail of a given distribution far more vigorously than does reference-dependent utility. A related implication is that reference-dependent utility may also tend to give more credit for gains above the

Figure 1. Power Utility vs. Reference-Dependent Utility

Notes: $Utility(\text{Gain}) = (\text{Wealth}/\text{Target})^{0.44} - 1$.
 $Utility(\text{Loss}) = 4.5[(\text{Wealth}/\text{Target})^{0.88} - 1]$.

target. Of course, these patterns are just tendencies, albeit ones that are associated with some common parameter values. The extent to which these tendencies apply will depend on the parameters adopted.

Figure 2 sheds further light on how the two utility functions evaluate a common distribution of outcomes. The figure plots the density of utility for each function under a lognormal distribution. To form the

Figure 2. Utility Densities of Outcomes under a Lognormal Distribution

series, a distribution of wealth outcomes is proposed that extends from 0.1 to 3.5 in intervals of 0.05. This distribution represents the range of wealth outcomes that might arise when investing over a horizon spanning many years. I then assumed that the range of outcomes was generated by a lognormal distribution⁸ and that the log of terminal wealth has a mean of 0 and a standard deviation of 0.5.⁹ These assumptions support estimation of the probability of observing a wealth outcome within each 0.05 interval. For instance, the probability of observing a wealth outcome between 0.50 and 0.55 (a 45%–50% wealth decline) is estimated at 3.31%; the probability of observing an outcome between 1.95 and 2.00 (a near doubling of wealth) is 0.80%.

I then calculated utility scores for each interval under both the power utility and reference-dependent utility functions shown in Figure 1. Finally, I multiplied the utility score for each interval by its probability, providing an estimate of the percentage contribution of each interval to expected utility under each utility function for a lognormal distribution.

Figure 2 confirms that power utility places far more weight on outcomes in the lower tail than outcomes in the upper part of the distribution. The reference-dependent utility function places weight on losses in a relatively more even fashion, with less skew toward weighting the lower tail. The reference-dependent function also gives greater credit for gains relative to the target. Finally, the reference-dependent utility

function places little weight on outcomes near the target wealth of 1.00. Figure 2 also reports certainty equivalent values of expected wealth, which stand at 0.72 for power utility and 0.88 for reference-dependent utility. This difference confirms that a power utility investor would require considerably more compensation for risk of loss than a reference-dependent investor would, at least for these chosen parameters.

Figure 1 and Figure 2 highlight the manner in which different utility functions may place different values on equivalent wealth outcomes and thereby arrive at quite different conclusions for the same distribution. These differences matter most when candidate portfolios differ in exposure to equities, which offer higher returns than, say, fixed-income securities but a wider distribution of outcomes. Indeed, reference-dependent utility functions may favor equities to the extent that they place relatively less penalty on the lower tail and give more credit for the upper part of the distribution. This aspect emerges when the interaction with time horizon is considered, as I discuss next.

Interaction with Horizon. When selecting a utility function, it helps to understand how different functions have differing implications for optimal portfolios as the investment horizon changes. Reference-dependent utility functions tend to give rise to an increasing preference for higher-returning assets, such as equities, as the horizon lengthens despite the higher volatility of

such assets. This effect does not occur under power utility, at least when the return or wealth distribution is independent and identically distributed (iid). A well-known characteristic is that optimal portfolios vary with horizon when the distribution is not iid, such as in the presence of serial correlation or when expected return and risk are time varying (Samuelson 1994; Barberis 2000; Campbell and Viceira 2005; Kritzman 2015). The discussion that follows abstracts from this element by assuming iid returns.

Whether investors should hold more equities over the long term is a hotly debated topic. The debate often occurs under the banner of “time diversification” (for an overview, see Thorley 1995; Bennyhoff 2009; Kritzman 2015). The theoretical roots, however, can be traced to discussions of the merits of “Kelly” investing, which date back to authors such as Kelly (1956), Latané (1959), Markowitz (1976), Samuelson (1971), and Merton and Samuelson (1974). The Kelly investing view has been interpreted as implying that assets with higher geometric returns should be increasingly favored as the horizon lengthens because such assets come to “almost stochastically dominate” assets with lower geometric returns (Leshno and Levy 2002). Nevertheless, some states typically remain under which the higher-returning asset with higher variance delivers lower wealth (Levy 2009). As a consequence, preferences play a central role in where one stands in this debate—as pointed out by Ziemba (2015) and Levy and Levy (2017). Preferences determine the relative weight placed on tail losses for the more volatile asset and thus whether equities might be preferred as the horizon lengthens.

Samuelson (1971) and Merton and Samuelson (1974) are credited with arguing that optimal portfolio weights do not change with the horizon. Their case is based on the assumption of CRRA preferences (power utility) as well as iid returns. Their results relate to the variance of wealth growing proportionately with time. Other authors have pointed out that constant weights need not hold in the case of a reference-dependent utility function (Benartzi and Thaler 1995; Bierman 1998; Levy and Levy 2017). The reason is that the probability of shortfall versus the target evolves with the horizon, depending, in part, on how the expected compound return compares with the rate at which the target grows. The action of compounding returns “shifts” the overall level of the wealth distribution upward. The trajectory of this shift depends on the level of expected returns, specifically their geometric mean. The overall shift in the distribution

then interacts with the manner in which gains and losses are asymmetrically evaluated under reference-dependent utility, as was discussed earlier, to create an increasing preference for assets offering higher compound returns as the horizon lengthens.

Figure 3 illustrates these aspects. The analysis is based on drawing 14 years of real returns for equities

Figure 3. Optimal Weights and Probability of Shortfall as Horizon Increases

and bonds from a normal distribution, with a mean of 6% and a standard deviation of 16% for equities and a mean of 1% and a standard deviation of 8% for bonds. The mean of 1% for 10-year bonds approximately represents real yields at the time of analysis; the 6% mean for equities accords with an equity risk premium of 5%, which broadly reflects historical experience. Panel A of Figure 3 plots optimal equity weights at horizons of 1 year to 14 years under power utility and reference-dependent utility functions using the same parameters as Figure 1, with a 3% real return target assumed to establish the reference point (W^*). The optimal equity weight does not vary with horizon under power utility, apart from the change arising from some variation in the random draws. This case is the one presented by Samuelson (1971) and Merton and Samuelson (1974). In contrast, under reference-dependent utility, equity weights unambiguously increase with horizon, attaining 100% after 12 years. This case underpins the work of Benartzi and Thaler (1995), Bierman (1998), and Levy and Levy (2017).

Panel B of Figure 3 helps explain why equity weights increase with horizon under the reference-dependent function by plotting shortfall probabilities. The probability of equities delivering a shortfall versus the target drifts lower from 43% in Year 1 to 34% in Year 14. Meanwhile, the probability of bonds delivering a shortfall rises from 60% to 86%. Thus, as the horizon lengthens, the proportion of the distribution falling into the loss area decreases for equities whereas it increases for bonds. Furthermore, the probability of equities underperforming bonds declines, reaching 20% in Year 14. This result indicates that the portion of the equity distribution that provides a worse outcome than bonds decreases as the horizon lengthens. Extending the charts would eventually lead to the “almost stochastic dominance” of equities over bonds found by Leshno and Levy (2002).

What Panel B of Figure 3 does not reveal is a lower tail where equities do much worse than bonds. This region is where the relatively more extreme penalty imposed by power utility on the lower tail comes into play. This penalty is sufficiently high to generate the static equity weights under power utility seen in Panel A of Figure 3, notwithstanding that equities do better than bonds over the bulk of the outcome distribution. Meanwhile, the reference-dependent utility function not only places a more “measured” penalty on this lower tail; it also gives more credit to the increasing slice of the distribution that falls

above the target as the horizon lengthens. Hence, the different ways in which each function balances gains versus losses interact with the manner in which the distribution of outcomes evolves over time to produce quite different optimal portfolio weights as the horizon lengthens. The main takeaway is that the form of utility function can be important for portfolio formation across horizons.

Utility Function Selection and Parameterization

An appropriate utility function should faithfully describe the objectives and preferences of the investor. In this section, I discuss how a suitable utility function might be chosen. I start by highlighting the dangers in picking an “off-the-shelf” utility function from the literature for use in any specific investment context. This warning supports the key message that the chosen functional form and parameters should make sense in the context in which the utility function will be used. I then provide some guidance on how to approach the task of choosing a utility function.

Dangers in Using an Off-the-Shelf Utility Function. The literature contains various studies that aim to identify and parameterize utility functions. The functions are specific to the circumstances in which they were estimated and are potentially unsuitable for a particular investor. Estimation typically arises from experimental studies, sometimes using students as subjects. A classic example is the cumulative prospect theory of Tversky and Kahneman (1992). The discussion here will focus on the estimate from that study of the curvature parameter on losses of 0.88 under what is a reference-dependent utility function in difference form.¹⁰ Applying a curvature parameter of less than 1.0 is consistent with imposing a proportionately *lower* penalty on losses as they increase in magnitude. For instance, the curvature parameter in isolation converts a \$10,000 loss to a utility value of $-\$3,311.3$ ($= -\$10,000^{0.88}$), meaning that the difference is “diluted” to about 33% of the raw value. A loss of \$20,000 is converted to a utility value of $-\$6,094.1$, which is about 30% of the raw value. Thus, a loss of \$20,000 is treated as proportionately less painful than a loss of \$10,000. (The heavy lifting in terms of penalizing losses is done by the weighting parameter.) One might question whether a proportionately decreasing penalty on losses is appropriate in all circumstances.

The 0.88 curvature parameter derives from the finding by Tversky and Kahneman (1992) of what they called “diminishing sensitivity” to gains and losses as outcomes move away from the reference point. It is highly questionable whether diminishing sensitivity to losses in the case of simple choices made by students can be transferred to other contexts, such as an institutional investor aiming to use utility functions to identify optimal portfolios. In some situations, it may be appropriate for the penalty on losses to increase at an increasing rate as outcomes fall further below target because the investor suffers increasing pain. For example, an investor who relies on drawing an income to meet living expenses could become increasingly destitute as his or her available wealth falls below target. Such a situation implies a curvature parameter in the realm of losses *in excess of 1.0*.

To drive the point home about the need to closely evaluate any candidate utility function, **Figure 4** presents the implications of applying two parameterizations appearing in the literature to both the difference and ratio forms of the reference-dependent utility function, thus generating four functions. The first parameterization follows Tversky and Kahneman (1992); it includes curvature parameters of 0.88 on both gains and losses and a weighting parameter of 2.25 on losses (1.0 on gains). The second parameterization follows Blake et al. (2013). It includes a curvature parameter of 0.44 on gains, a curvature parameter of 0.88 on losses, and a weighting parameter of 4.5 on losses (1.0 on gains). Target wealth of \$20,000 is assumed. Panel A of Figure 4 plots utility values under the four functions for outcomes ranging from \$0 (i.e., -\$20,000 in difference form and 0.0 in ratio form) up to \$40,000 (respectively, +\$20,000 and 2.0). Utility values under the difference form appear on the left scale; those under the ratio form appear on the right scale. The scaling permits the implications for evaluating outcomes to be directly compared. Panel B of Figure 4 examines the relative scores placed on gains and losses under each function, by reporting the absolute score for a gain as a percentage of the absolute score for a loss of an equivalent magnitude. Numbers for power utility are also shown for comparison. The aim of Figure 4 is to highlight the implications of the choice of both functional form and parameters. The aim is not to comment on which combination of parameters is “right.”

Utility functions are used to rank prospects, so the units are not directly meaningful in isolation. The main attributes to consider are the broad shape of

the function and the implications for evaluating outcomes. Panel A of Figure 4 shows that the Tversky and Kahneman (1992) parameters impose a relatively modest penalty on losses compared to the discount applied to gains. Panel B of Figure 4 confirms this information by revealing that a gain is afforded a utility value with a magnitude of about 40% of the utility value of an equivalent loss.¹¹

The parameters selected by Blake et al. (2013) are much more loss averse than those of Tversky and Kahneman (1992). The differing implications of the same parameters under the two functional forms is notable. Under the difference form (used by Blake et al.), an extremely heavy discount is imposed on gains, which are afforded minimal credit. Panel B of Figure 4 reveals that the magnitude of the utility score on a gain is *less than 1%* of the score attached to an equivalent loss. The implication is that the investor is primarily concerned with minimizing the magnitude of losses versus the target and does not care much about the possibility of gains. Panel A of Figure 4 reveals that applying the parameters of Blake et al. in the ratio form leads to a heavier penalty on losses but also a less aggressive discount on gains. Panel B of Figure 4 reports that the ratio of utility scores for gains and losses of an equivalent magnitude is around 10% in this case. Panel B also highlights the heavy and increasing penalty imposed by power utility on outcomes in the lower part of the distribution.

The key message is that adopting a utility function and associated parameters simply because they appear in the academic literature may not be an effective approach. A utility function chosen in this way could provide a poor representation of the objectives and preferences of the investor. The better path is to forfeit the comfort of being able to cite existing literature for validation and aim to identify the utility function that is appropriate in the circumstances.

Choosing a Utility Function. Choosing a utility function is a large topic worthy of a separate paper. Some general guidance on how the task might be approached is offered here. The process is illustrated later in the descriptions of selecting utility functions for four representative investors.

The starting point is to form an understanding of the investor’s objectives and preferences over outcomes. These findings are then reflected in a utility function that can be used to evaluate those outcomes. Key aspects to consider in choosing a utility function

Figure 4. Parameterizations in the Literature under Reference-Dependent Utility

A. Utility Values under the Four Functions for a Range of Outcomes

B. Ratio of Utility Values for Gains vs. Losses

Value		Difference		Ratio		Power Utility (CRRA = 4)
		Tversky & Kahneman (1992)	Blake et al. (2013)	Tversky & Kahneman (1992)	Blake et al. (2013)	
+\$2,000 vs. -\$2,000	±10%	44.4%	0.8%	43.9%	10.7%	54.8%
+\$4,000 vs. -\$4,000	±20%	44.4	0.6	43.4	10.4	29.6
+\$8,000 vs. -\$8,000	±40%	44.4	0.4	42.3	9.8	7.9
+\$16,000 vs. -\$16,000	±80%	44.4	0.3	39.8	8.7	0.01

Notes: Panel B reports the absolute value of the ratio of the utility values for gains and losses of equivalent magnitude, relative to target wealth of \$20,000. The counterpart ratios reported for power utility are based on equivalent values for wealth in absolute terms; for example, +\$4,000 vs. -\$4,000 is the ratio of power utility values for wealth of \$24,000 vs. \$16,000.

are listed below. The aim is to initially identify the functional form and then specify the parameters. Determining the parameters might be aided by using some of the techniques presented in earlier sections—for example, plotting the function and

analyzing the relative scores placed on various outcomes across the feasible range.

- **Outcomes over which utility is to be defined.** The first step is to identify what outcomes the investor cares about, over which utility should

be defined: withdrawals from the portfolio (e.g., income or consumption), end-period wealth (i.e., portfolio value), or some combination of the two.

- *Whether the investor is concerned with raw levels or a target.* Concern with the level of withdrawals and/or wealth implies a function that directly evaluates outcomes, such as power utility and other forms of the HARA class. If a target exists, then some kind of reference-dependent utility function may be appropriate.
- *Nature of any reference point.* Reference points may take a variety of forms, including target withdrawals and/or wealth, return targets (e.g., a consumer price index plus some additional percentage), performance delivered by a benchmark or peers, funding ratios, and capital ratios. An additional issue is whether the reference point is time varying—for example, resetting as some form of habit persistence. The nature of any target, if one exists, needs to be captured in the utility function.
- *Investment horizon.* The horizon over which utility is to be measured can be influential, as discussed previously. Utility should be defined with respect to the time period of concern to the investor. The evaluation horizon may differ from the period over which wealth will be deployed. For instance, investors with defined-contribution retirement plans might be characterized quite differently depending on whether they are concerned with their near-term balance or with the potential retirement income they can draw over the long run.
- *Time preference.* When a stream of outcomes is generated over time, the question arises whether the investor has a preference for attaining those outcomes earlier rather than later. If so, a time preference parameter may need to be incorporated into the utility function.
- *Views about the value attached to the residual balance.* In some circumstances, the value attributed to any residual balance may have distinctive attributes that need to be captured in the utility function. A classic example is a bequest motive, where the value placed on bequests may differ from that placed on withdrawals or wealth accruing to the investors themselves.
- *Preferences for differing outcomes.* Parameterizing the utility function sets the value that the investor will place on outcomes that are lower or higher within the distribution. A helpful approach is to categorize preference as scores

assigned to “good” versus “bad” outcomes rather than a matter of risk aversion. Understanding these preferences and encoding them in the utility function is perhaps the most challenging aspect of the analysis. As discussed previously, the key consideration is the relative scores placed on outcomes across the range. For instance, extreme aversion to being in the lower tail would suggest using a high CRRA under power utility or relatively large curvature and/or weighting parameters on losses under reference-dependent utility, probably in conjunction with a small curvature parameter on gains. Another issue is the existence of a *point of ruin*. One example is where a retiree is left destitute if her or his income is too low or the retirement balance falls to zero. Another example is where erosion of capital leads to bankruptcy for an organization. In such situations, the utility function should substantially ramp up the penalty on losses as the ruin point approaches.¹² Under this case, evaluations become dominated by avoiding ruin as the first priority.

Various techniques can assist in teasing out investor objectives and preferences in order to select and parameterize a utility function. Potential methods include observing behaviors (revealed preferences), choice experiments, surveys or questionnaires, and open discussions at the board level or with the client.

An Approach for Forming Portfolios

An approach is set out here for identifying optimal portfolios by using utility functions. The approach will be illustrated in the next section. It entails four steps:

1. Choose a utility function.
2. Specify a plan for withdrawal and evaluation.
3. Project joint distributions of accumulated wealth outcomes for the assets.
4. Solve for optimal portfolio weights.

Step 1: Choose a Utility Function. The task of selecting and parameterizing a utility function was discussed earlier, so only a few additional points are necessary here. Although any utility function could be used, the three forms of utility function described in the previous section will suffice in many circumstances. They have the flexibility to capture the objectives and preferences of many classes of investor and have received considerable recognition in the

academic literature. They also offer the advantage of being time additive, which simplifies the aggregation of utility over time.

Step 2: Specify a Plan for Withdrawal and Evaluation.

The approach requires specifying withdrawals over the investment horizon as well as the combination of withdrawals and terminal portfolio value to be evaluated by the chosen utility function. Withdrawals constitute distributions from the portfolio for consumption or redeployment. They must be specified in their entirety; the sequence of portfolio values will depend on any prior withdrawals. The terminal portfolio value is the residual available at the end of the horizon. Terminal portfolio value need not be intended for consumption or redeployment but, nevertheless, should be evaluated to the extent that residual wealth is considered of importance to the investor. The utility function must be used to evaluate any outcomes that are conditional on investment performance and, hence, asset weights. Withdrawals that are fixed amounts may be excluded from the evaluation plan for the purpose of optimization because their utility value will be a constant. Including fixed withdrawals in the calculations does no harm, however, and may be useful if expected utility estimates are of interest in themselves—for example, for estimating a certainty equivalent.

Although a withdrawal plan may be specified in many ways, defining withdrawals as a percentage of portfolio value can be convenient from an analytical perspective because it ensures that portfolio value remains positive. A related approach would be to follow Waring and Siegel (2015) by estimating withdrawals as payments arising from a virtual annuity that is recalculated conditional on asset values and interest rates within the simulated paths.¹³ Defining withdrawals as fixed amounts may lead to wealth falling to zero in certain circumstances. This aspect complicates the estimation of utility, which will hinge on the evaluation of *either* any residual left after meeting the planned fixed withdrawals *or* the utility consequences of the portfolio value declining to zero before all planned withdrawals are satisfied. This concern does not invalidate the approach, but specifying fixed withdrawals requires careful consideration of how any residual portfolio value and/or zero withdrawals will be evaluated under the utility function.¹⁴

The example plans for withdrawal and evaluation will be briefly described here to convey a sense of what is entailed, with more detail appearing later in

the article. The *private investor* is primarily concerned with wealth at the end of the horizon. Therefore, she evaluates terminal portfolio value only and does not allow for withdrawals. The perpetual *endowment fund* evaluates withdrawals on the basis of a predetermined percentage of portfolio value up to Year 10 plus the terminal portfolio value. The terminal value is considered representative of the remaining pot of wealth to satisfy withdrawals beyond Year 10. The *defined-benefit pension fund* faces fixed withdrawals over 15 years, but the sponsor desires to avoid deficits lasting more than 3 years. Thus, the sponsor evaluates the funding ratio at Year 3. Withdrawals in the interim are included in the withdrawal plan but excluded from the evaluation plan because they are fixed. Finally, the *retired individual* is concerned with income over the 30 years he expects to remain alive. Thus, this individual evaluates the entire stream of withdrawals over that period.

Two other situations may be instructive. A possible variation for a retired individual would be to allow for a bequest motive by including the terminal portfolio value in the evaluation plan, perhaps by imposing a discount on the bequest by means of a scaling factor. Bell et al. (2017b), following Lockwood (2014) in the treatment of residual portfolio value, demonstrated this variation for power utility. Under reference-dependent utility, the bequest might be evaluated by setting a target terminal portfolio value of zero, so that the bequest is treated as a gain, and perhaps imposing a discount by adjusting the curvature parameter or applying a weighting parameter of less than 1.0. For *investment managers*, a suitable plan might involve focusing on terminal portfolio value at the end of the horizon over which the manager expects to be evaluated by investors. For instance, if the investors are willing to afford the manager a three-year grace period, portfolio value might be evaluated at the end of Year 3. The fact that most mandates are based on benchmarks suggests a reference-dependent utility function, with a target that reflects projected wealth from investing in the benchmark. If the benchmark is an index, the target might refer to an index fund or exchange-traded fund.¹⁵ If the benchmark is a return target, the target becomes the terminal portfolio value consistent with achieving the target return. The target might be varied conditionally within each simulated path, providing a connection between the manager's performance and the market state. For instance, a real return target of the consumer price index plus 3% would be modeled by reference to projected inflation within each path.

Step 3: Project Joint Distributions of Accumulated Wealth Outcomes for the Assets.

The approach is based on modeling assets in terms of the accumulated wealth they deliver and aggregating these estimates into a portfolio value from which any withdrawals are to be subtracted. Modeling accumulated wealth is particularly helpful when projecting long-term outcomes over many periods. Potential methods for modeling the joint distribution of accumulated wealth arising from investing in a group of candidate assets include the following:

- Simulating from historical data by jointly drawing observations from the same period for each asset
- Applying statistical models, such as vector autoregressive or regime-switching models, which are specified in a way that accounts for any interrelationships between assets
- Applying structural models that impose relationships between key variables and the assets—for example, the stochastic model of Wilkie (1984), which is widely used in actuarial practice
- Using value-based models with input variables shared among assets; for example, the economy jointly drives the return on equity that is input into a plowback equity model and the bond yield that is input into the bond model
- Using scenario analysis (see Gosling 2010)

The modeling process should satisfy two requirements. First, covariance between any assets should be embedded in the joint distribution of wealth outcomes. In mean-variance analysis, this requirement is achieved by specifying a covariance matrix. For multiperiod analysis over long periods, the modeling may be structured by simulating a series of joint paths or specifying a range of scenarios, which act as “states” to which each asset responds. Second, wealth generated from reinvestment should be accounted for appropriately. Reinvestment of the cash flows generated by an asset accounts for an increasing portion of wealth as the investment horizon lengthens. It is thus important for long-term investors to ensure that reinvestment is properly modeled. The common approach is to form accumulation indexes that assume distributed cash flows are reinvested in the asset itself. In some circumstances, the appropriate approach may be to model the reinvestment process explicitly. One example is using a plowback or similar model for equities, thus accounting for the marginal

return on reinvestment. Another example is explicitly modeling the reinvestment of coupons and principal in fixed-income analysis.

For the illustrations in a later section, I used two basic methods to simulate accumulated wealth for equities and 10-year bonds, which suffice to illustrate the approach. The modeling assumed that dividends and coupons were reinvested in the asset itself. Joint distributions were generated in two ways—through random draws and from historical data:

- *Random draws.* A series of 10,000 paths of accumulated wealth was formed by randomly drawing real returns from a normal distribution, with allowance for uncertainty over the expected equity return. For equities, the expected return of 6% was treated as a random variable and modeled by drawing different expected returns for each path from a normal distribution with a mean of 6% and a standard error of 2%. The standard deviation for yearly equity returns is 16%. For 10-year bonds, expected returns have a known mean of 1% and a standard deviation of 8%. Apart from the random equity mean, these assumptions are the same as those underpinning Figure 3. Allowing for uncertainty over the expected equity return has the effect, relative to the iid case, of causing the distribution of expected wealth to fan out as the horizon increases. The randomly drawn data assume that equity and bond returns are uncorrelated.
- *Historical data.* Yearly data for US equity returns, long bond yields, and inflation from 1871 to 2016 were taken from Robert Shiller’s website.¹⁶ The data were used to create a series of (overlapping) return paths over time spans of up to 10 years. To create a series of 10-year bond returns, I assumed that a 10-year bond paying yearly coupons equal to the reported long bond yield was purchased at the end of each year. This bond was notionally sold as a 9-year bond at the end of the year.¹⁷ The sale proceeds plus the coupon were then reinvested in another 10-year bond. The accumulated wealth series were formed in nominal terms. They were deflated by the inflation index during each path to generate a real wealth series and, hence, a real return series. This method implicitly embeds the historical covariance between equities and bonds as well as any serial correlation in returns.

The method generated wealth paths representing a total of 135 actual 10-year episodes that

were experienced by investors. The mean real return (standard deviation) for the sample is 8.0% (18.0%) for equities and 2.4% (7.9%) for bonds, and the correlation between the two assets is 0.163 (0.045 in nominal terms). Another feature of the data is that bonds have a positive serial correlation of 0.155, whereas equity correlations are close to zero, at 0.012.

Step 4: Solve for Optimal Portfolio

Weights. Forming optimal portfolios requires generating a set of paths, each comprising a sequence of portfolio values and associated withdrawals over the period. The sequence is formed by conditioning within each path on the wealth arising from investing in each asset, the withdrawal plan, and the asset weights. Closing portfolio values for each year are calculated by adjusting opening portfolio values for the change in wealth arising from investing at candidate asset weights and then deducting withdrawals. Deducting withdrawals implicitly accounts for any sequencing effects. The utility function is used to evaluate withdrawals and the terminal portfolio value for each path in accordance with the evaluation plan, resulting in a utility score for each path. These scores are averaged across paths to form an expected utility measure. Finally, the optimal portfolio is identified by finding the asset weights that maximize expected utility. Further details of the modeling process, including formulas, now follow.

The initial aim is to estimate the accumulated portfolio value (PV) that is available to support withdrawals (D) for a forecast period of $t = h$ within each simulated path p (i.e., state). Equation 4 is used to estimate the sequence of portfolio values in a forward direction, so that portfolio values incorporate withdrawals up until the beginning of period t . Equation 5 provides estimates of the withdrawal for period t determined as a percentage of the pre-withdrawal portfolio value. The handling of fixed withdrawals was discussed in step 2 and would amount to simply deducting a predetermined amount for $D_{p,t}$:

$$PV_{p,t} = PV_{p,t-1} \left[1 + x_e R_{e,p,t} + (1 - x_e) R_{b,p,t} \right] - D_{p,t}, \quad (4)$$

and when $D_{p,t}$ is a function of $PV_{p,t}$,

$$D_{p,t} = c_t PV_{p,t-1} \left[1 + x_e R_{e,p,t} + (1 - x_e) R_{b,p,t} \right], \quad (5)$$

where

$PV_{p,t}$ = portfolio value in path p for period t

x_e = weight in equities

$R_{e,p,t}$ = equity return in path p for period t

$R_{b,p,t}$ = bond return in path p for period t

c_t = percentage of portfolio value to be withdrawn in period t

$D_{p,t}$ = withdrawals in path p for period t

Withdrawals in each path and terminal portfolio value are scored by using the chosen utility function and then aggregated to derive the total utility score associated with that path. This approach is straightforward when the withdrawal is equal to the terminal portfolio value at the end of the horizon, which gives rise to a single utility score for that path. When a stream of withdrawals and possibly the terminal portfolio value are being evaluated, aggregating utility scores across time becomes necessary.

When utility is time additive, the standard approach of simply summing discounted utility can be applied through Equation 6:

$$U_{(.),p} = \sum_{t=1}^{t=h} \tau^t U_{(.)}(D_{p,t}) + \tau^h U_{(.)}(PV_{p,h}), \quad (6)$$

where

$U_{(.),p}$ = aggregate utility for path p under (time-additive) utility function $U_{(.)}$

τ = time preference parameter

Equation 6 allows for a time preference parameter (τ), which may be required when the investor has a preference for early withdrawals ($\tau < 1.0$). Time preference is ignored by setting $\tau = 1.0$. This assumption is made in the analysis in the next section, which is conducted in real terms.

The assumption that utility may be summed across time within each path is usually valid for power utility and for reference-dependent utility in its difference form, where the utility estimates represent a direct transformation of dollars versus target. It may not be valid for reference-dependent utility functions in their ratio form because the scaling by target has the potential to undermine the comparability of utility scores over time when similar percentage deviations

from target represent different “value” to the investor. For instance, a value of $W/W^* = 1.10$ when W equals \$11,000 and W^* equals \$10,000 may not be equivalent to the value when W equals \$1,100 and W^* equals \$1,000. A key example is where outcomes comprise a mix of withdrawals and terminal portfolio value with different value implications. In such circumstances, some sort of scaling may be required. One approach is to specify weights to permit the utility estimates to be aggregated in a meaningful manner:¹⁸

$$U_{RDU(R),p} = \sum_{t=1}^{t=h} y_t U_{RDU(R),t}, \quad (7)$$

where

$$U_{RDU(R),p} = \text{aggregate reference-dependent utility in ratio form for path } p$$

$$y_t = \text{weight placed on utility generated in period } t$$

The weights y_t should be chosen in accordance with the circumstances. If the withdrawal for each period is of equal importance, then simple addition will suffice. In many instances, weighting by the respective target outcomes (i.e., $W_t^* / \sum_{t=1}^{t=h} W_t^*$) or perhaps the projected outcomes (i.e., $W_t / \sum_{t=1}^{t=h} W_t$) may be appropriate. If earlier withdrawals are considered more relevant to the investor than later withdrawals, this time preference can be accommodated by further adjusting the weights for a discount factor. In any case, the key aim is to aggregate the utility score in a manner that encapsulates the preferences of the investor. Judgment is arguably more important than mathematical precision.

The utility estimates for each path are then weighted by probability to generate expected utility:

$$E[U] = Pr_p \sum_{p=1}^n U_{(.),p}, \quad (8)$$

where

$$U_{(.),p} = \text{utility in path } p, \text{ estimated by either Equation 6 or Equation 7}$$

$$Pr_p = \text{probability of path } p$$

In the analysis in the next section, expected utility is estimated by taking a simple average of utility

across all paths, on the assumption that each path is equally likely.

Finally, the portfolio weights that maximize the value of Equation 8 are found. In the two-asset example reflected in Equation 4 and used in the next section, this step amounts to solving for the optimal equity weight, x_e .¹⁹

Illustrations for Four Investor Types

The choice and use of utility functions in forming portfolios is illustrated here by applying the approach to four representative investors. The aim is to demonstrate the process and draw out some key themes.

Private Investor. The private investor is investing a “pot of wealth” she wants to grow over time, with no specific return or wealth target. She is averse to large losses that would be sustained through to the end of her investment horizon but is not overly concerned about volatility along the path. This situation implies the use of power utility to evaluate the terminal portfolio value at the end of the investment horizon. I use this as an opportunity to demonstrate how optimal portfolios may vary with horizon and risk aversion. The approach is to generate results for two horizons, 3 years and 10 years, for CRRA parameters ranging from 1 (log utility) to 8. Accumulated wealth associated with investing in equities and 10-year bonds is simulated using historical data from Shiller’s website. **Exhibit 1** summarizes the setting, outlines the key inputs for steps 1 through 3, and reports the optimal equity weights under step 4. Equity weights are plotted in **Figure 5**.

The optimal equity weights range from 32.5% up to 100%, depending on risk aversion and horizon, with a median of 65%. For this dataset, equity weights are approximately 15% higher for the 10-year horizon than for the 3-year horizon. This difference reflects serial correlation in the data, specifically the fact that real bond returns are persistent (serial correlation of 0.155 versus 0.012 for equities), which increases their risk relative to equities over long periods.²⁰ The plausibility of these results is revealed by comparing the weights to observed portfolios held by private investors. The 65% median is close to the average equity weight of 62.4% reported by the American Association of Individual Investors for December 2018.²¹ The spread of my estimates also seems in the ballpark. The 2013 Survey of Consumer Finances reports on the asset allocation of households in their initial defined-contribution plans.²² Panis and Brien

Exhibit 1. Private Investor Setting

Step	Input/Output		
1. Utility function	<ul style="list-style-type: none">Power utilityCRRA from 1 (log utility) to 8 investigatedAggregation of utility across time not required because analysis focused on terminal portfolio value		
2. Plan for withdrawal and evaluation	<ul style="list-style-type: none">Evaluation of terminal portfolio valueTwo horizons investigated: 3 years and 10 years		
3. Assets and projection of wealth outcomes	<ul style="list-style-type: none">Equities and 10-year bondsSimulations using 10-year historical real returns based on Shiller data, with the 10-year bond rolled over each year		
4. Portfolio formation: optimal portfolios	<ul style="list-style-type: none">Equity weights		
	CRRA	3-Year Horizon	10-Year Horizon
	1	100.0%	100.0%
	2	100.0	100.0
	3	83.4	98.1
	4	63.7	79.1
	5	51.4	66.6
	6	43.1	57.6
	7	37.1	50.9
	8	32.5	45.6

Figure 5. Optimal Portfolio Weights for the Private Investor with Various Inputs

(2016) analyzed these data and found that 26% opt for equities only, 11% opt for interest-earning assets only, and 62% opt for a mix. The survey results for IRAs/Keogh accounts showed that 34% opt for equities only and 52% opt for a mix.

Endowment Fund. The endowment fund is investing a pool of capital (the “corpus”) intended to support philanthropic activities in perpetuity. All potential stakeholders—the board, donors, and recipients—have agreed on withdrawals (distributions) equal to 5% of the portfolio at the end of each year and an objective of maintaining, if not increasing, the real value of withdrawals over the long term. The situation calls for evaluating the entire stream of future withdrawals with the use of a reference-dependent utility function in perpetuity. A practical way of analyzing this situation is to project and evaluate withdrawals plus terminal portfolio value over some long forecast horizon, with the terminal portfolio value taken to represent the residual pot of wealth available to support philanthropic activities beyond the horizon. A horizon of 10 years is considered sufficiently long and accommodates the effective use of the historical Shiller data in simulating the asset returns and, hence, the accumulated wealth paths.²³ The target is based on maintaining both withdrawals over the forecast period and the terminal portfolio value in real terms. Because of the differing nature of the withdrawals (yearly amounts) and the terminal portfolio value (the larger pot of wealth), the reference-dependent utility function in its ratio form is a natural choice. The resulting utility values are aggregated through weighting by the real value of the target components. The resulting

weights are 3.33% on each of the 10 withdrawals and 66.7% on the terminal portfolio value.

The parameterization of the utility function draws on the assumption that the yearly withdrawal becomes a budget to be allocated without any firm spending commitments. This situation reduces aversion to losses by limiting the adverse implications of a reduction in portfolio value and, therefore, withdrawals. The endowment thus has latitude to pursue higher returns and the resulting greater capacity to support philanthropic activities, albeit at some risk of diminishing the corpus. The utility function is thus parameterized for moderate aversion to losses versus the target and to give reasonable credit for gains.

Figure 6 plots the utility function for the chosen parameters, which include curvature parameters of 0.50 on both gains and losses and a weighting function of 3.0 on losses. For example, this curve implies a utility value of +0.12 for exceeding the target by +25% and -0.40 for a shortfall of -25%. The curve also applies an increasing rate of penalty on losses and discount on gains—with, for example, utility values of +0.22 for exceeding the target by +50% and -0.88 for a shortfall of -50%.

Exhibit 2 summarizes the endowment fund setting. The optimal portfolio weight is 100% in equities under this setup. Basically, the analysis recommends that the endowment fund run an aggressive strategy in pursuit of high long-term returns to increase the potential to “do more good” over time. This result stems from applying a reference-dependent utility function with moderate loss aversion over a long horizon.

Figure 6. Utility Function for the Endowment Fund

Exhibit 2. Endowment Fund Setting

Step	Input/Output
1. Utility function	<ul style="list-style-type: none"> Reference-dependent utility, ratio form Parameters: <ul style="list-style-type: none"> Curvature on gains vs. target: 0.50 Curvature on losses vs. target: 0.50 Weighting on gains vs. target: 1.0 Weighting on losses vs. target: 3.0 Time preference: 1.0 (no discounting) Target: Maintenance of the real value of withdrawals and portfolio value Aggregation: Utility values weighted by target (implies weights of 3.33% on withdrawals and 67.7% on terminal portfolio value)
2. Plan for withdrawal and evaluation	<ul style="list-style-type: none"> Yearly withdrawals equal to 5% of portfolio value Evaluation of withdrawals from Year 1 to Year 10 plus terminal portfolio value at Year 10 (the latter represents the wealth available to support withdrawals beyond Year 10)
3. Assets and projection of wealth outcomes	<ul style="list-style-type: none"> Equities and 10-year bonds Simulations using 10-year historical real returns based on Shiller data, with the 10-year bond rolled over each year
4. Portfolio formation: optimal portfolios	<ul style="list-style-type: none"> Equity weight: 100%

The outcome could change if more conservative parameters were applied or the horizon were reduced. For instance, halving the curvature parameter on gains to 0.25, doubling the curvature parameter on losses to 1.0, and imposing a weighting parameter of 6.0 on losses generates an optimal equity weight of slightly under 90%. If these parameters were combined with an evaluation of withdrawals plus terminal portfolio value over a three-year horizon, the optimal equity weight would decline to 53%. A shortening of the horizon might be justified if stakeholders were sensitive to maintaining the value of the corpus at every point in time or if the endowment fund had no flexibility to vary the assumed 5% distribution rate and withdrawals needed to be maintained to support spending commitments. In any event, the analysis needs to be tailored to the situation of the investor.

The equity weights discussed here broadly accord with the relatively aggressive asset mixes held by actual US endowments. NACUBO (2018) reported that US college and university endowments held only 12% in fixed income and cash in 2017; the

residual was in equities and, largely, growth alternatives (although 19% was in marketable alternatives, including hedge funds).

Defined-Benefit Fund (Liability-Driven Investing). The defined-benefit (DB) fund is investing to satisfy a stream of pension payments that are fixed in nominal terms. The analysis is undertaken from the perspective of the fund sponsor, which is concerned about the reported funding ratio. The sponsor has some tolerance for funding deficits if sustained for only a limited period but would much prefer that deficits not persist beyond three years. The sponsor is also keen to avoid making additional contributions to top up the fund assets but places some value on surpluses as a buffer and anticipates benefiting from any residual value after all obligations are met. The situation implies that the funding ratio should be evaluated in Year 3 by using a reference-dependent utility function in ratio form that is parameterized for relatively high risk aversion in relation to deficits, while attaching a relatively modest value to surpluses. This evaluation

is achieved by applying the parameters of Blake et al. (2013), which include a curvature parameter of 0.44 on surpluses, a curvature parameter of 0.88 on deficits, and a weighting parameter of 4.5 on deficits. A feel for these parameters was provided earlier, in Figure 4, which showed that they deliver utility values for gains in the order of 10% of those for losses of equivalent magnitude. This information implies a much stronger preference for avoiding deficits than generating surpluses.

The pension payments and their valuation are modeled in a simple way here, albeit sufficient to draw out the key concepts. The liability is a fixed nominal pension payment (i.e., term annuity) spanning 15 years. It is valued by discounting at the prevailing 10-year bond rate.²⁴ The Shiller data are used to generate a series of simulated paths for bond yields, asset returns, and accumulated wealth over a 3-year horizon. To value the liability, the remaining 12 pension payments at the end of Year 3 are discounted at the prevailing 10-year bond yield for that path. This step establishes a path-specific value for the target—that is, $W_{p,t=3}^*$. The setup embeds a strong correlation between the liability and the 10-year bond, so the bond plays the role of the low-risk hedging asset for the liability. At the same time, equities offer a higher return but are mismatched with the liability. Evaluating the funding ratio at the end of Year 3 bypasses the withdrawals made in Year 1 through Year 3, which are fixed and make no difference to the optimization. These withdrawals, nevertheless, influence the terminal portfolio value at Year 3, which is reduced by the value of withdrawals under the sequential modeling. **Exhibit 3** summarizes the setting, **Figure 7** plots the utility function, and **Figure 8** plots the optimal equity weights.

The analysis is conducted for a range of initial funding ratios to demonstrate how the optimal equity weight is sensitive to this ratio, as evident in Figure 8. The mechanism is as follows: When the fund commences at a funding ratio near 1.00, the optimal portfolio becomes heavily weighted in bonds in an attempt to secure the funding ratio. This reflects the hedge that bonds provide for the liability coupled with high aversion to deficits. As the initial funding ratio moves increasingly into deficit, equities are more and more favored because they increase the probability of reaching the target funding ratio of 1.00 by Year 3. At the same time, bonds become unattractive because they only lock in the existing deficit. As the initial funding ratio moves further into surplus, the probability of falling back into deficit

by Year 3 falls and scope emerges to add equities to capture some of the upside. This pattern should be familiar to those who have ventured into the analysis of liability-driven investing. Note that the U shape in Figure 8 arises because of the assumption that the sponsor places value on surpluses. If the analysis were rerun with zero utility value placed on surpluses, the main change would be that equity weights would fall to zero at an initial funding ratio of 1.09 and beyond.

Comparing these results with weightings observed in the field is difficult because of the simplified setting of this illustration. Nevertheless, the pattern broadly accords with the finding by Panis and Brien (2015) of a negative relationship between funding ratios and the weight in aggressive assets for a sample of US DB funds, where funding ratios were mainly below 1.25. The median aggressive asset weight for their sample was 63%, which is consistent with the optimal equity weight at a funding ratio of 85%–90% in the setting examined here.

Retired Individual. Retirement savings in this section are analyzed for an individual in a simplified Australian setting, where well-defined values are available for contribution and withdrawal rates as well as target income. The individual has just retired at 65 years of age and has transferred his superannuation into a pension account. (The next section extends the analysis to a comparison of alternative strategies in a life-cycle setting.) He intends to draw on this account over a 30-year period, until he turns 95, when he expects to die. His target is to draw annual income of at least A\$24,506 in real terms over the 30 years, which reflects the “modest” retirement standard estimated by the Association of Superannuation Funds of Australia (ASFA) in September 2017. His withdrawal plan is to follow the minimum drawdown requirements set by the Australian government, which start at 5% at age 66 and rise to 11% for ages 91–95 (see Figure 10 for further details). No allowance is made in this analysis for the many other issues in retirees’ lives, including whether this individual is eligible for the age pension, the existence of other assets outside his pension account, uncertainty over his longevity, and any bequest motive. Isolating the analysis from these influences focuses the discussion on how the choice of utility function might affect asset weights without too many confounding effects. Given that the age pension provides a risk-free hedge against both investment losses and longevity risk, its exclusion is an important deviation from practice.

Exhibit 3. DB Fund Setting

Step	Input/Output																		
1. Utility function	<ul style="list-style-type: none"> Reference-dependent utility, ratio form <i>Parameters</i> (following Blake et al. 2013): <ul style="list-style-type: none"> Curvature on gains vs. target: 0.44 Curvature on losses vs. target: 0.88 Weighting on gains vs. target: 1.0 Weighting on losses vs. target: 4.5 <i>Target</i>: Funding ratio of 1.00 <i>Aggregation</i>: Not required because analysis focuses on terminal funding ratio at end of Year 3 																		
2. Plan for withdrawal and evaluation	<ul style="list-style-type: none"> Evaluation of funding ratio at Year 3, calculated as portfolio value divided by present value of liabilities estimated by discounting at the 10-year bond rate in Year 3 in each path Withdrawals in Year 1 through Year 3 are fixed. Although not evaluated by the utility function, they affect portfolio value at the end of Year 3. 																		
3. Assets and projection of wealth outcomes	<ul style="list-style-type: none"> Equities and 10-year bonds Simulations using 3-year historical nominal returns based on Shiller data, with the 10-year bond rolled over each year 																		
4. Portfolio formation: optimal portfolios	<ul style="list-style-type: none"> Optimal portfolio depends on initial funding ratio: <table> <tr> <th>Initial Funding Ratio</th><th>Equity Weight</th></tr> <tr> <td>0.70</td><td>100.0%</td></tr> <tr> <td>0.80</td><td>91.9</td></tr> <tr> <td>0.90</td><td>51.5</td></tr> <tr> <td>1.00</td><td>15.8</td></tr> <tr> <td>1.10</td><td>26.3</td></tr> <tr> <td>1.20</td><td>51.3</td></tr> <tr> <td>1.30</td><td>69.4</td></tr> <tr> <td>1.40</td><td>82.9</td></tr> </table>	Initial Funding Ratio	Equity Weight	0.70	100.0%	0.80	91.9	0.90	51.5	1.00	15.8	1.10	26.3	1.20	51.3	1.30	69.4	1.40	82.9
Initial Funding Ratio	Equity Weight																		
0.70	100.0%																		
0.80	91.9																		
0.90	51.5																		
1.00	15.8																		
1.10	26.3																		
1.20	51.3																		
1.30	69.4																		
1.40	82.9																		

The situation suggests evaluating the entire stream of withdrawals by using a reference-dependent utility function in difference form and ignoring any terminal portfolio value.²⁵ It also implies imposing a high penalty on any shortfall versus target income, given that the retiree is relying on withdrawals from his pension account for living expenses. The opportunity is taken to contrast the results for two utility function parameterizations, one deemed “preferred” and the other reflecting the parameters used by

Blake et al. (2013) in a UK context. **Exhibit 4** summarizes the setting, **Figure 9** plots the two utility functions, and **Figure 10** plots the optimal equity weights.

Relative to Blake et al. (2013), the preferred parameters are designed to give more credit for exceeding the target while imposing a penalty on losses that increases in relative magnitude with the shortfall. More credit is given for exceeding the target by

Figure 7. Utility Function for DB Fund

Figure 8. Optimal Equity Weight for DB Fund

assuming a curvature parameter on gains of 0.80, versus 0.44 in Blake et al., under the assumption that the retiree places some value on the possibility of a better life style. The increasing penalty on shortfall versus the target is achieved by assuming a curvature parameter on losses in excess of 1 (specifically, 1.1)

instead of the 0.88 assumed by Blake et al. To ensure that the penalty on losses is not overly large, it is coupled with a weighting parameter on losses of 2.0 (versus the 4.5 assumed by Blake et al.). The analysis was conducted for a range of starting balances, which turns out to be an important variable.

Exhibit 4. Retired Individual Setting

Step	Input/Output																																				
1. Utility function	<ul style="list-style-type: none">Reference-dependent utility, difference form<i>Parameters</i>—two sets examined: <table><tr><th>Parameter</th><th>Preferred</th><th>Blake et al. (2013)</th></tr><tr><td>Curvature on gains</td><td>0.80</td><td>0.44</td></tr><tr><td>Curvature on losses</td><td>1.10</td><td>0.88</td></tr><tr><td>Weighting on gains</td><td>1.0</td><td>1.0</td></tr><tr><td>Weighting on losses</td><td>2.0</td><td>4.5</td></tr></table> <ul style="list-style-type: none"><i>Target:</i> Annual withdrawal of \$24,506 maintained in real terms; equivalent to ASFA standard for a modest life style estimated at September 2017<i>Aggregation:</i> Utility summed across all 30 years, with time preference of 1.0 (no discounting)	Parameter	Preferred	Blake et al. (2013)	Curvature on gains	0.80	0.44	Curvature on losses	1.10	0.88	Weighting on gains	1.0	1.0	Weighting on losses	2.0	4.5																					
Parameter	Preferred	Blake et al. (2013)																																			
Curvature on gains	0.80	0.44																																			
Curvature on losses	1.10	0.88																																			
Weighting on gains	1.0	1.0																																			
Weighting on losses	2.0	4.5																																			
2. Plan for withdrawal and evaluation	<ul style="list-style-type: none">Evaluation of withdrawals over 30 years between age 66 and age 95, reflecting the minimum drawdown rate for Australian retirees: <table><tr><th>Age</th><th>Drawdown</th></tr><tr><td>65–75</td><td>5%</td></tr><tr><td>76–80</td><td>6</td></tr><tr><td>81–85</td><td>7</td></tr><tr><td>86–90</td><td>9</td></tr><tr><td>91–95</td><td>11</td></tr></table> <ul style="list-style-type: none">Terminal portfolio value not evaluated (i.e., no bequest motive)	Age	Drawdown	65–75	5%	76–80	6	81–85	7	86–90	9	91–95	11																								
Age	Drawdown																																				
65–75	5%																																				
76–80	6																																				
81–85	7																																				
86–90	9																																				
91–95	11																																				
3. Assets and projection of wealth outcomes	<ul style="list-style-type: none">Equities and 10-year bondsSimulations using 10,000 draws of real returns from lognormal distributions, including real equity returns with a random mean of 6% (standard error of 2%) and a standard deviation of 18% and real bond returns with a mean of 1% and a standard deviation of 8%																																				
4. Portfolio formation: optimal portfolios	<ul style="list-style-type: none">Optimal portfolio depends on initial balance and preference parameters (see Figure 10 for plot of full series of optimal weights): <table><tr><th>Selected Balances</th><th>Preferred Parameters</th><th>Blake et al. Parameters</th></tr><tr><td>\$400,000 and below</td><td>100.0%</td><td>100%</td></tr><tr><td>\$600,000</td><td>75</td><td>69</td></tr><tr><td>\$800,000</td><td>69</td><td>56</td></tr><tr><td>\$1.0 million</td><td>80</td><td>49</td></tr><tr><td>\$1.2 million</td><td>100</td><td>46</td></tr><tr><td>\$1.7 million (Blake et al. min. weight)</td><td>100</td><td>44</td></tr><tr><td>\$2.0 million</td><td>100</td><td>47</td></tr><tr><td>\$3.0 million</td><td>100</td><td>68</td></tr><tr><td>\$4.0 million</td><td>100</td><td>83</td></tr><tr><td>\$5.0 million</td><td>100</td><td>95</td></tr><tr><td>\$5.5 million (Blake et al. max. weight)</td><td>100</td><td>100</td></tr></table>	Selected Balances	Preferred Parameters	Blake et al. Parameters	\$400,000 and below	100.0%	100%	\$600,000	75	69	\$800,000	69	56	\$1.0 million	80	49	\$1.2 million	100	46	\$1.7 million (Blake et al. min. weight)	100	44	\$2.0 million	100	47	\$3.0 million	100	68	\$4.0 million	100	83	\$5.0 million	100	95	\$5.5 million (Blake et al. max. weight)	100	100
Selected Balances	Preferred Parameters	Blake et al. Parameters																																			
\$400,000 and below	100.0%	100%																																			
\$600,000	75	69																																			
\$800,000	69	56																																			
\$1.0 million	80	49																																			
\$1.2 million	100	46																																			
\$1.7 million (Blake et al. min. weight)	100	44																																			
\$2.0 million	100	47																																			
\$3.0 million	100	68																																			
\$4.0 million	100	83																																			
\$5.0 million	100	95																																			
\$5.5 million (Blake et al. max. weight)	100	100																																			

Figure 9. Utility Functions for Retired Individual

Figure 10. Optimal Equity Weight for Retired Individual

The analysis is based on random draws of 10,000 paths for accumulated wealth with uncertainty over the expected equity returns, as detailed earlier.

Figure 10 reveals a U-shaped pattern with respect to the starting balance, with the shape of the curve depending on the utility function parameters. Both the preferred and Blake et al. (2013) parameters lead to 100% equity at balances of \$400,000 and below. Equities are favored because their higher returns increase the odds of attaining the target. Optimal

equity weights diverge considerably beyond that point. Under the preferred parameters, the optimal equity weight declines to 69% at a starting balance of \$800,000 and then returns to 100% at a balance of \$1.2 million. Under the parameters of Blake et al., the equity weight declines more sharply and bottoms at just below 44% at a balance of \$1.7 million. It then gradually increases. The chart only extends to \$2.5 million, but the equity weight reaches 100% at an (unlikely) balance of \$5.5 million.

The utility curves plotted in Figure 9 explain the weighting patterns seen in Figure 10. The Blake et al. (2013) parameters afford little credit for exceeding the target, meaning that the optimization is dominated by outcomes in the realm of losses. The result is a much more extended U shape. A reluctance emerges to increase equity exposure until the balance is high enough to render only a small probability of falling into the loss region. The preferred parameters give more credit than those of Blake et al. to the possibility of exceeding the target, resulting in a greater propensity to take on equity risk. The Blake et al. parameters might be appropriate for a retired individual with extreme aversion to shortfall who does not care much about exceeding the target. The preferred parameters might be more suitable for a retired individual who is willing to balance the downside against the upside.²⁶

Alternative Strategies for Retirement Savings.

I now apply utility-based analysis to compare retirement saving strategies over an individual's life cycle. Although only a few investment strategies are considered here, this style of analysis could be applied more broadly. For instance, it could be used to examine whether changing the contribution or withdrawal strategy might be beneficial. Indeed, any aspect that can be expressed in terms of its impact on withdrawals or portfolio value may be evaluated in terms of whether it enhances utility. Comparing alternatives can be useful when the optimal solution is practically out of reach or when we want to evaluate a proposed marginal change versus current practice.

The simulation analysis of the prior section is extended to incorporate contributions and investment performance over an accumulation phase starting at age 25. Contributions equaling 9.5% of wage income are assumed, in line with the current mandated rate in Australia. Wage income is set at male average weekly ordinary-time earnings as of November 2017, for an annual income of \$86,580, growing at 1% per year in real terms until retirement. Utility is defined over withdrawals (i.e., income) during retirement in line with the minimum drawdown rules as noted in the previous section. Estimates are provided for the two parameterizations of the reference-dependent utility function and for power utility with CRRAs of 2, 5, and 8. Five investment strategies are applied over the 70-year life cycle: 100% bonds, an investment rule under which the equity weight is set at "100 – age" (see Dahlquist, Setty, and Vestman 2018), a target date fund, a balanced portfolio with constant 60% equities and 40% bonds, and 100% equities. The target date fund is modeled on the Vanguard Target Retirement Fund series, which holds 90% in equities until 25 years from retirement, then transitions to 30% equities until 7 years after retirement. The Vanguard glide path is similar to the average in US (see Morningstar 2018) and Australian life-cycle funds (see Chant, Mohankumar, and Warren 2014), although the Australian life-cycle funds distribute the balance at the point of retirement.

Table 1 reports selected estimates arising from 10,000 simulations. Utility estimates are converted into certainty equivalent income (CEI), which is the constant real income stream over the 30 years of

Table 1. Comparison of Investment Strategies Applied over the Life Cycle

Strategy	Certainty Equivalent Income (CEI)					Median Retirement Income	5th Percentile Income	% of Periods Income Is below Target
	Reference-Dependent Utility		Power Utility					
	Preferred	Blake et al. (2013)	CRR A 2	CRR A 5	CRR A 8			
100% bonds	15,763	16,379	13,193	7,894	4,877	15,773	6,104	83
100 – age	22,285	22,614	26,611	13,669	6,738	32,245	11,818	32
Target date fund	23,166	23,125	31,563	13,131	5,557	40,565	12,698	24
Constant 60/40	23,464	23,306	33,131	10,047	3,509	43,439	13,105	21
100% equities	40,151	23,279	36,439	3,016	963	71,804	10,899	17

Note: The strategy generating the highest CEI for each utility function is in boldface.

retirement that generates the same expected utility as that arising under a risky strategy.²⁷ Also reported are the median income during retirement, the fifth percentile income as a guide to the lower tail, and the percentage of times the strategy fails to generate the income target under the reference-dependent utility function. The strategies are ordered by the median income they generate under the simulations.

The analysis is instructive in various ways. First, it illustrates how the preferred strategy varies with utility function. The CEI estimates suggest that the preferred strategy is 100% equities for investors with low loss or risk aversion, a constant 60/40 balanced strategy under reference-dependent utility with the Blake et al. (2013) parameters, and a 100 – age strategy under power utility with a high CRRA of 5–8. This highlights how applying an inappropriate utility function can result in an unsuitable strategy being selected.

Second, the analysis generates estimates of the relative value of the strategies to an investor and, by implication, the cost of applying an inappropriate utility function. Consider a conservative investor who might be best represented by power utility with a CRRA of 5. Assume such an investor were in a 60/40 balanced fund, based on, say, analysis that used a reference-dependent utility function with the parameters of Blake et al. (2013). This investor could improve his situation by more than \$3,000 a year as measured by CEI through switching to a target date fund or the 100 – age strategy, both of which take some risk off the table as the investor ages. If the individual were currently 100% invested in fixed income, switching to one of these strategies would offer a potential gain of more than \$5,000.

Finally, I have demonstrated how utility-based analysis can rank alternatives by using a single measure that takes into account the entire distribution of outcomes. Arguably, this process is more elegant than attempting to trade off, say, the median income against the risk metrics—such as income at the fifth percentile or the percentage of periods that the strategy fails to reach the target.

Static vs. Dynamic Analysis

The approach presented in this study offers a method for forming portfolios by using utility functions that can be readily applied by practitioners who have a good knowledge of such spreadsheet packages as Microsoft Excel and who want to avoid

complex mathematical or numerical methods. Here the main technique for reducing complexity is shutting down the dynamics. This is done by specifying *deterministic* withdrawal plans and solving for static asset weights as if they will be maintained throughout the forecast period. This approach precludes the possibility that investors might *anticipate* changing their future investment and/or withdrawal strategies. This structure accommodates a simpler analysis and modeling in a forward direction.

Extending to *dynamic* analysis requires considering what dynamic elements might be included and whether they would make any substantial difference to the results. Dynamics can be incorporated on a number of levels.

The most advanced and potentially complex method assumes that both the investment and withdrawal strategies will both be *reoptimized* in response to future developments. Such problems require full-scale dynamic programming that solves for the optimal asset weights and withdrawals, typically as a recursive problem (i.e., working backwards).²⁸ These methods apply either closed-form mathematical solutions, which often constrain the analysis to make it tractable (for instance, see Soupé, Heckel, and Leote de Carvalho 2016), or the application of complex numerical techniques (for example, see Butt and Khemka 2015; Khemka and Butt 2017). Cochrane (2014) suggested that few investors understand such models, let alone use them.²⁹

A less complex approach entails specifying predetermined rules for how asset weights and/or withdrawals (and, possibly, contributions) might respond to investment returns along the path. These specifications will not yield a fully optimal dynamic solution but can support an analysis with a relatively simple structure. For example, rules may be set for how asset weights or withdrawals or both respond after realizations of low or high asset returns or for how they respond as a function of distance from the target.³⁰ One issue with applying predetermined rules is that it may become necessary to iterate joint sets of investment and withdrawal strategies to converge on a preferred solution that is not totally ad hoc. This task can be cumbersome if a large “grid” of possible strategy combinations needs to be evaluated.

Another simplification technique is to limit the withdrawal plan to a single or a few obvious alternatives and focus on identifying an “acceptable” dynamic investment strategy for the specified plan(s). The modeling task is thereby reduced to locating the

set of dynamic asset weights that maximizes utility for a candidate plan, with the aim of identifying the preferred combination of investment strategy and withdrawal plan from a manageable number of iterations. This method may be suitable when finding the optimal withdrawal plan is not the key focus and it is acceptable to propose and evaluate a limited number of withdrawal strategies.

In any event, the question is whether dynamic analysis would lead to a substantially different asset allocation *today* or whether the static weights arising from applying the basic approach might suffice. The answer hinges on the extent to which the optimal weights for today depend on what could happen tomorrow. The importance of this issue will vary with the circumstances. In many instances, to conduct the analysis on the basis of what is known today may suffice and the possible response to future developments can be ignored. This approach does not imply that the investor intends to stick with a static strategy regardless of what happens. The analysis can be rerun and the strategy readjusted periodically. Cochrane (2014) implied that this approach is what most practitioners currently use in any event.

To gauge the difference that dynamic analysis might make, I applied stochastic dynamic programming to identify the optimal investment strategy for the DB fund introduced earlier.³¹ The analysis draws on the techniques used by Butt and Khemka (2015) and Khemka and Butt (2017). It involves applying recursive analysis to create a grid of optimal equity weights for each period conditional on the realized funding ratio (which is treated as the “state” variable) and the distribution of future expected asset returns.³² In this analysis, I again estimated optimal equity weights at the beginning of Year 1 for initial funding ratios in the range of 0.60 to 1.40. Utility estimates were also collected. Under the dynamic analysis, the Year 1 equity weight was 3.1% greater, on average, than under the previous static analysis. Although the equity weight differences were uneven across the range of initial funding ratios,³³ most differences were relatively modest in magnitude and a similar U-shaped equity weight function emerged. Thus, the representative DB fund would not substantially alter its current equity weights knowing it has the option to revise its future equity weights. Further, the utility gains averaged +0.028 (median +0.019),³⁴ which indicates that the option to revise weights is of some modest value. Importantly, the tenor of the results did not change in any major way under the dynamic analysis.

One faces a trade-off between reducing complexity and accepting what could be a suboptimal solution. Whether some level of dynamic analysis is needed to properly address the problem at hand will depend on the circumstances. In particular, if jointly solving for the optimal investment and withdrawal strategies is important, then full-scale dynamic optimization may be necessary. If the withdrawal plan can be largely predetermined, then the approach presented in this article may suffice (or it may be extended to accommodate dynamics to a limited degree). Perhaps fortuitously, applying dynamic analysis to the DB fund led to only modest changes. Although generalizing from specific cases is dangerous, these results hint that the approach in its basic form may deliver acceptable results across a range of situations.

Conclusion

Utility-based analysis can provide an effective and flexible means of identifying optimal portfolios in practical settings. The investment industry has tended to rely on mean-variance analysis, in part because of tradition and familiarity with the approach, which is widely taught in universities and colleges. Mean-variance analysis is essentially a single-period approach and hence may not accord with the objectives and preferences of many investors. For instance, it may not be effective if the horizon is long and involves a stream of withdrawals over time or if the investor is concerned with a target outcome. This study demonstrated how utility-based analysis offers an alternative to mean-variance optimization that can be tailored to a wide variety of investors with different objectives and preferences, while generating plausible results. I presented an approach in which the analysis can be conducted entirely in Microsoft Excel without having to resort to dynamic optimization, albeit by imposing some constraints on the scope of the analysis.

My key contention is that utility functions should be selected and parameterized with a view to tailoring to the investor. Researchers can have a natural desire to rely on findings from the academic literature in order to seek validation. The problem is that the utility functions and parameter estimates appearing in the academic literature could be quite unsuitable for the investor in question. A much better approach is to design the utility function and related analysis to reflect the circumstances. Going through the process of understanding the investor’s objectives and preferences and encoding them in a utility function not

only leads to more meaningful outputs than using an off-the-shelf approach but also invites the practitioner to address the question of what really matters to the investor along the way.

Editor's Note

Submitted 1 October 2018

Accepted 20 March 2019 by Stephen J. Brown

Notes

1. A *certainty equivalent* is a guaranteed value that an investor would accept now rather than taking a chance on a higher, but uncertain, value in the future.
2. Merton (1971) noted that HARA utility is a rich class that accommodates absolute and relative risk aversion that can be increasing, decreasing, or constant in wealth. It includes the power, quadratic, exponential, and log utility functions. The utility function attributed to Epstein and Zin (1989) and Weil (1989) decouples the “elasticity of intertemporal substitution” from risk aversion, allowing any preference for earlier resolution of uncertainty to be evaluated separately from aversion to volatility. Although this function is widely used in the academic literature, the intuition behind this form of utility can be hard to grasp, and the function needs to be solved recursively, which adds to complexity.
3. Log utility can be seen as a limiting case of power utility as CRRA approaches 1. Rubinstein (1976) argued that log utility has many attractive features from a modeling perspective. For many investors, however, it will place an insufficient penalty on low outcomes.
4. The first derivative of Equation 1 with respect to W is W^{-CRRA} .
5. Prospect theory entails a broader framework than the value function, including a preliminary stage where prospects are “edited” down to a subset for evaluation and application of decision weights that transform the probabilities attached to outcomes. It proposes an S-shaped function that implies risk aversion for gains but risk seeking for losses.
6. These utility functions can be treated as subcases of the general form of reference-dependent utility function presented here.
7. The general form of the HARA family of utility functions entails three parameters and provides flexibility to manipulate absolute and relative risk aversion (see Merton 1971, p. 389).
8. This assumption equates to compounding returns over multiple periods, with $\ln(1 + \text{Return}_t)$ being independent and normally distributed.
9. This standard deviation would arise from investing over 10 years with a yearly standard deviation of 15.8%; that is, $0.50 = [(0.158 \times 0.158) \times 10]^{0.5}$.
10. The curvature parameter on losses of 0.88 has been adopted by such authors as Benartzi and Thaler (1995), Blake et al. (2013), and Levy (2016). Levy (2016) also investigated values of between 0.5 and 1.0 for this parameter. He commented, “In practice the loss aversion is presumably more profound than what is estimated in laboratory experiments” (p. 1421).
11. The ratio is exactly 44.4% in the difference form because the same curvature parameter is applied to gains and losses, while the weighting parameter applies a linear transformation. The ratio is nonlinear under the ratio form of the function because taking a ratio imposes a nonlinear transformation on the outcomes before the utility function is applied.
12. The literature is often concerned that utility functions be continuous and differentiable. Although these characteristics are required for analytical solutions, they are not necessary in numerical approaches that simply assign and aggregate scores attached to outcomes. Utility might even be specified as a schedule of scores, rather than an explicit function, when numerical methods are being used.
13. The approach outlined in this study can provide a way of setting the allocation between riskless and risky assets in the Waring–Siegel (2015) method by including their liability-hedging TIPS (Treasury Inflation-Protected Securities) ladder as one of the candidate assets.
14. A zero withdrawal becomes undefined (infinite) under power utility. This problem might be handled by specifying a minimum utility value—for example, by substituting the utility value of \$1 for any withdrawal less than \$1. In reference-dependent utility, consideration might be given to whether the maximum negative utility value achieved at \$0 imposes an adequate utility penalty for running out of money. In any event, the aim is to ensure that an appropriate utility value be attached to states where the portfolio value and, hence, withdrawals decline to zero.
15. Another approach would be to conduct the analysis in benchmark-relative terms.
16. Available at www.econ.yale.edu/~shiller/data.htm.
17. The reported long bond yield for that period was used in pricing the bond.
18. Another approach, similar to that described for bequests, is to apply a discount to terminal portfolio value directly within the utility function.
19. If the problem is sufficiently well behaved, the analysis might be conducted in Microsoft Excel by using the Solver add-in to find the optimal weights.
20. The positive serial correlation for bonds arises from their greater exposure to inflation, which tends to be persistent.
21. See www.aaii.com/assetallocationsurvey for the latest figure.
22. The report is available at www.federalreserve.gov/econres/scf_2013.htm.

23. The 10-year horizon is a compromise. Extending the length of the data blocks drawn from the Shiller data reduces sample size and would not change the result.
24. The pension payments are set at a value of \$0.094 per \$1 of liability, so that the discounted value of the payments is \$1 at the bond yield assumed to prevail in period $t = 0$ of 4.92%. This yield represents the average long bond yield in the Shiller dataset over the period of analysis. This assumption establishes a baseline that dovetails with the historical data.
25. If the retiree placed a value on any bequest, the terminal portfolio value would also need to be evaluated. Bell et al. (2017a) provided an example of such an application.
26. The findings in this section are hard to square with investor behavior for a range of reasons, one of which is that this analysis reflects a highly simplified setting. Also, data on the strategies followed by retirees are limited, and decumulation strategies are still in the development stage. Indeed, utility-based analysis might be a tool to help develop better retirement strategies than we currently observe in practice.
27. CEI is estimated as a function of expected utility by rearranging the utility function. In the case of power utility, for instance, CEI, age 66–95
$$= \left\{ \left[E(\text{Utility})/30 \right] \times (1 - \text{CRRA}) \right\}^{1/(1 - \text{CRRA})}.$$
28. Note that utility functions such as those of Epstein and Zin (1989) and Weil (1989) require recursive solutions.
29. Admitting dynamic portfolio management may give rise to positions aimed at hedging changes in the investment opportunity set, in the sense of Merton (1973).
30. The analysis in the prior section assumes rebalancing back to target asset weights, which amounts to a simple rule of this type.
31. Recall that the withdrawals were fixed for this investor.
32. Initially, a grid of optimal equity weights for Year 3 is established by optimizing utility at the end of Year 3 conditional on the Year 2 funding ratio. The next step is to create a similar grid of equity weights for Year 2 conditional on the realized Year 1 funding ratio, again optimizing utility at the end of Year 3 and assuming that the Year 3 equity weights will follow the Year 3 grid. The final step is to estimate the optimal equity weights at the beginning of Year 1 to optimize utility at the end of Year 3, under the assumption that the equity weights for Year 2 and Year 3 will subsequently be updated in line with both grids.
33. The Year 1 equity weight differences between the dynamic and static strategies were mainly positive and ranged up to about +8%, although the optimal dynamic equity weights were modestly lower for initial funding ratios in the 0.90–1.00 range.
34. Under the reference-dependent utility function, this change might be roughly interpreted as a risk-adjusted change in wealth.

References

- Adler, Timothy, and Mark Kritzman. 2007. "Mean-Variance versus Full-Scale Optimisation: In and Out of Sample." *Journal of Asset Management* 7 (5): 302–11.
- Anthonisz, Sean A. 2012. "Asset Pricing with Partial-Moments." *Journal of Banking & Finance* 36 (7): 2122–35.
- Barberis, Nicholas. 2000. "Investing for the Long Run When Returns Are Predictable." *Journal of Finance* 55 (1): 225–64.
- Bell, David, Estelle Liu, and Adam Shao. 2017a. "Member's Default Utility Function for Default Fund Design Version 1 ('MDUF v1'): Design." Technical Paper No. 1 (21 March). http://membersdefaultutilityfunction.com.au/wp-content/uploads/2017/06/technical_paper_1_mdof_v1___derivation.pdf.
- . 2017b. "Member's Default Utility Function for Default Fund Design Version 1 ('MDUF v1'): Static Models." Technical Paper No. 2 (21 March). http://membersdefaultutilityfunction.com.au/wp-content/uploads/2017/06/technical_paper_2_mdof_v1___static_models.pdf.
- Benartzi, Shlomo, and Richard H. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics* 110 (1): 73–92.
- Bennyhoff, Donald G. 2009. "Time Diversification and Horizon-Based Asset Allocations." *Journal of Investing* 18 (1): 45–52.
- Bierman, Harold, Jr. 1998. "A Utility Approach to the Portfolio Allocation Decision and the Investment Horizon." *Journal of Portfolio Management* 2 (1): 81–87.
- Blake, David, Douglas Wright, and Yumeng Zhang. 2013. "Target-Driven Investing: Optimal Investment Strategies in Defined Contribution Pension Plans under Loss Aversion." *Journal of Economic Dynamics and Control* 37 (1): 195–209.
- Butt, Adam, and Gaurav Khemka. 2015. "The Effect of Objective Formulation on Retirement Decision Making." *Insurance: Mathematics and Economics* 64 (September): 385–95.
- Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–51.
- Campbell, John Y., and Luis M. Viceira. 2005. "The Term Structure of the Risk-Return Trade-Off." *Financial Analysts Journal* 61 (1): 34–44.
- Chant, Warren, Mano Mohankumar, and Geoff Warren. 2014. "MySuper: A New Landscape for Default Superannuation Funds." CIFS Paper No. 020/2014.
- Cochrane, John H. 2014. "A Mean-Variance Benchmark for Intertemporal Portfolio Theory." *Journal of Finance* 69 (1): 1–49.
- Dahlquist, Magnus, Ofer Setty, and Roine Vestman. 2018. "On the Asset Allocation of a Default Pension Fund." *Journal of Finance* 73 (4): 1893–936.
- Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4): 937–69.

- Estrada, Javier, and Mark Kritzman. 2018. "Evaluating Retirement Strategies: A Utility-Based Approach." Working paper (March). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=313512.
- Fishburn, Peter C. 1977. "Mean-Risk Analysis with Risk Associated with Below-Target Returns." *American Economic Review* 67 (2): 116–26.
- Goetzmann, William, Jonathan Ingersoll, Matthew Spiegel, and Ivo Welch. 2007. "Portfolio Performance Manipulation and Manipulation-Proof Performance Measures." *Review of Financial Studies* 20 (5): 1503–46.
- Gosling, Susan. 2010. "A Scenarios Approach to Asset Allocation." *Journal of Portfolio Management* 37 (1): 53–66.
- Grishchenko, Olesya V. 2010. "Internal vs. External Habit Formation: The Relative Importance for Asset Pricing." *Journal of Economics and Business* 62 (3): 176–94.
- Hogan, William W., and James M. Warren. 1972. "Computation of the Efficient Boundary in the ES Portfolio Selection Model." *Journal of Financial and Quantitative Analysis* 7 (4): 1881–96.
- Holthausen, Duncan M. 1981. "A Risk-Return Model with Risk and Return Measured as Deviations from a Target Return." *American Economic Review* 71 (1): 182–88.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–92.
- Kelly, John L., Jr. 1956. "A New Interpretation of Information Rate." *Bell System Technical Journal* 35 (4): 917–26.
- Khemka, Gaurav, and Adam Butt. 2017. "Non-Parametric Integral Estimation Using Data Clustering in Stochastic Dynamic Programming: An Introduction Using Lifetime Modelling." *Risks* 5 (4): 57.
- Kingston, Geoff, and Susan Thorp. 2005. "Annuitization and Asset Allocation with HARA Utility." *Journal of Pension Economics and Finance* 4 (3): 225–48.
- Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-Dependent Risk Attitudes." *American Economic Review* 97 (4): 1047–73.
- Kritzman, Mark. 2015. "What Practitioners Need to Know ... About Time Diversification (corrected)." *Financial Analysts Journal* 71 (1): 29–34.
- Latané, Henry A. 1959. "Criteria for Choice among Risky Ventures." *Journal of Political Economy* 67 (2): 144–55.
- Leshno, Moshe, and Haim Levy. 2002. "Preferred by 'All' and Preferred by 'Most' Decision Makers: Almost Stochastic Dominance." *Management Science* 48 (8): 1074–85.
- Levy, Haim. 2016. "Aging Population, Retirement, and Risk Taking." *Management Science* 62 (5): 1415–30.
- Levy, Haim, and Moshe Levy. 2017. "Prospect Theory and Investment Horizon." Working paper, Hebrew University of Jerusalem (22 October). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3056969.
- Levy, Moshe. 2009. "Almost Stochastic Dominance and Stocks for the Long Run." *European Journal of Operational Research* 194 (1): 250–57.
- Lockwood, Lee M. 2014. "Incidental Bequests and the Choice to Self-Insure Late-Life Risks." NBER Working Paper 20745 (December).
- Markowitz, Harry. 1976. "Investment for the Long Run: New Evidence for an Old Rule." *Journal of Finance* 31 (5): 1273–86.
- Merton, Robert C. 1971. "Optimal Portfolio and Consumption Rules in a Continuous-Time Model." *Journal of Economic Theory* 3 (4): 373–413.
- . 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41 (5): 867–87.
- Merton, Robert C., and Paul A. Samuelson. 1974. "Fallacy of the Log-Normal Approximation to Optimal Portfolio Decision-Making over Many Periods." *Journal of Financial Economics* 1 (1): 67–94.
- Morningstar. 2018. "Vanguard Target Retirement Target-Date Fund Series Report." <https://news.morningstar.com/pdfs/STUSA04OVV.pdf>.
- NACUBO. 2018. "2017 NACUBO-Commonfund Study of Endowments." National Association of College and University Business Officers.
- Panis, Constantijn W.A., and Michael J. Brien. 2015. "Asset Allocation of Defined Benefit Pension Plans." Advanced Analytical Consulting Group and Deloitte (19 November). <https://www.dol.gov/sites/default/files/ebsa/researchers/analysis/retirement/asset-allocation-of-defined-benefit-pension-plans.pdf>.
- . 2016. "Financial Asset Holdings of Households in the United States in 2013." Advanced Analytical Consulting Group and Deloitte (21 January). <https://www.dol.gov/sites/default/files/ebsa/researchers/analysis/retirement/financial-asset-holdings-of-households-in-the-us-in-2013.pdf>.
- Rubinstein, Mark. 1976. "The Strong Case for the Generalized Logarithmic Utility Model as the Premier Model of Financial Markets." *Journal of Finance* 31 (2): 551–71.
- Samuelson, Paul A. 1971. "The 'Fallacy' of Maximizing the Geometric Mean in Long Sequences of Investing or Gambling." *Proceedings of the National Academy of Sciences of the United States of America* 68 (10): 2493–96.
- . 1994. "The Long-Term Case for Equities." *Journal of Portfolio Management* 21 (1): 15–24.
- Soupé, François, Thomas Heckel, and Raul Leote de Carvalho. 2016. "Portfolio Insurance with Adaptive Protection." *Journal of Investment Strategies* 5 (3): 1–15.
- Starmer, Chris. 2000. "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk." *Journal of Economic Literature* 38 (2): 332–82.
- Stutzer, Michael. 2003. "Portfolio Choice with Endogenous Utility: A Large Deviations Approach." *Journal of Econometrics* 116 (1–2): 365–86.
- Tarlie, Martin B. 2017. "Investment Horizon and Portfolio Selection." Working paper (8 March). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2854336.
- Thorley, Steven R. 1995. "The Time Diversification Controversy." *Financial Analysts Journal* 51 (3): 68–76.
- Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.

Unser, Matthias. 2000. "Lower Partial Moments as Measures of Perceived Risk: An Experimental Study." *Journal of Economic Psychology* 21 (3): 253–80.

Veld, Chris, and Yulia V. Veld-Merkoulova. 2008. "The Risk Perceptions of Individual Investors." *Journal of Economic Psychology* 29 (2): 226–52.

Waring, M. Barton, and Laurence B. Siegel. 2015. "The Only Spending Rule Article You Will Ever Need." *Financial Analysts Journal* 71 (1): 91–107.

Weil, Philippe. 1989. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *Journal of Monetary Economics* 24 (3): 401–21.

Wilkie, A. David. 1984. "A Stochastic Investment Model for Actuarial Use." *Transactions of the Faculty of Actuaries* 39 (January): 341–403.

Ziemba, William T. 2015. "A Response to Professor Paul A. Samuelson's Objections to Kelly Capital Growth Investing." *Journal of Portfolio Management* 42 (1): 153–67.

RESEARCH FOUNDATION REVIEW 2018

The *Research Foundation Review 2018* summarizes the offerings from the CFA Institute Research Foundation over the past year—books, literature reviews, workshop presentations, and other relevant material.

For more information or to download a free copy, visit cfa.is/review2018

© 2019 CFA Institute. All rights reserved.