

RETIREMENT MANAGEMENT JOURNAL

A reprinted article from Volume 10, Number 1, 2021

Efficiency Metrics

MAXIMIZING UTILITY OF WITHDRAWALS IN RETIREMENT AND THE
EFFICIENCY OF REQUIRED MINIMUM DISTRIBUTIONS

By Chester Chambers, PhD



INVESTMENTS & WEALTH INSTITUTE®

Efficiency Metrics

MAXIMIZING UTILITY OF WITHDRAWALS IN RETIREMENT AND THE EFFICIENCY OF REQUIRED MINIMUM DISTRIBUTIONS

By Chester Chambers, PhD

ABSTRACT

We focus on a simplified problem for a risk-averse retiree seeking to maximize utility associated with annual spending and a remaining value at the end of the problem horizon when the funds are extracted from a portfolio that includes a risk-free and a risky asset. To organize discussions about this setting we use a novel metric that we label “efficiency.” This measurement compares the utility derived from annual withdrawals and the final value with a benchmark based upon the maximum sustainable spending rate that could have been chosen under perfect information. We use 10,000 scenarios developed using historical data to evaluate efficiency over a family of spending rules and asset allocations. In the process, we develop insights about the structure of such rules and asset allocation schemes. The metric used here can be useful to planners and advisors who wish to present a comparison of policies built around a single value. Some withdrawal strategies will outperform the benchmark for many investors. In particular, we find that strategies built around required minimum distribution (RMD) rules are surprisingly efficient and also have characteristics that mitigate against some effects of market risk, longevity risk, and sequencing risk. Our findings suggest that advisors and clients will be well served by comparing policies in a consistent way and that we all should more seriously consider withdrawal strategies built around RMD rules.

INTRODUCTION

Like many people, I was thrust into a role that included managing the liquidation of assets during the last years of my father’s retirement. On the surface this was a very easy problem to handle. He received a monthly pension check as a result of 20 years of military service. Due to this service his healthcare costs were essentially zero. He received a separate monthly check as a result of another 20 years of work as a civil servant. In addition, he received a small Social Security check. Consequently, his regular income was more than sufficient to cover all essential expenses. Despite my strong protests, he never would consider a lifestyle that involved living a penny above these means. However, he did have a small nest egg of savings bonds and cash that he would consider “playing with” if I could explain, to his satisfaction, how it would be managed.

My intent was to give him license to spend some and give some away as he saw fit. At the same time, he wanted to be absolutely sure that most of this money would be left for his grandchildren.

In dealing with this setting a number of factors were particularly prominent. First, under no circumstances would this portfolio be allowed to hit a value of zero. Second, a high level of complexity was not going to be tolerated. I was able to talk him into holding one risky asset, but I wasn’t going to win any arguments pushing that idea any further. Finally, risk aversion was high. Even though this money was divorced from day-to-day expenses, losses would be felt far more than any comparable gains. In thinking about ways to organize discussions about these issues, I settled on a notion of efficiency that would allow me to come up with a “score” for any policy that could be compared with any other policy to provide a simple rationale for one approach over another. This will never take the place of a full retirement and estate plan, but I suspect that many planners can benefit from consideration of the approach that was used.

For lots of reasons, William F. Sharpe often gets quoted as saying that retirement planning is “the nastiest, hardest problem in finance” (Ritholtz 2017). As virtually all readers of this journal will attest, there is at least a grain of truth in this sentiment. This is in large part because retirees experience a wide range of emotions that get in the way of rational decision-making. Retirees have longevity concerns, and avoiding outliving one’s money is almost universally accepted as the first priority. Simultaneously, most retirees have at least some interest in leaving a bequest to heirs or charities. In addition to these dominant interests, we also must account for risk aversion in the sense that uncertain payout levels are undesirable. Also, many advisors are all too familiar with the fact that performance always will be compared, consciously or not, with some idiosyncratic benchmark or with anecdotes from friends or peers. This human trait leads to a sense of regret if things do not work out as planned that is not realizable until after the passage of some period of time. Finally, we have the specter of RMDs, which typically is portrayed as the scourge of the earth because it leads to the universally hated payment of taxes.

With all this in mind, we would like to focus on a small idea that reflects consideration of these issues and presents a novel way to inform discussions about planning for similar settings. In particular, we employ a notion of efficiency that takes into account historical performance, extreme risk aversion, anticipation of regret, and the desire to bequest something after all is said and done. In essence, we use randomized data from the historical record, along with an asset allocation rule, some withdrawal policy, and a simple utility function that relates to both withdrawals and bequests to state the efficiency of a policy as a single value. This approach provides a quick and dirty way to sort a collection of plans and to organize discussions about the trade-offs involved. Along the way we touch upon three results that we found rather surprising. First, variability of payouts is not the end of the world, and it is far better to formally account for high levels of risk aversion than to speak about planning as though evolution of payout levels will never happen. Second, RMDs may not be nearly as bad as most people think, in the sense that rules built around these requirements are surprisingly efficient. Finally, fairly aggressive asset allocations are less risky than more conservative ones in that they are more efficient and are much more likely to reduce regret.

Let us contemplate a retiree with a portfolio of assets to be managed with the intent of maximizing the enjoyment or utility of annual withdrawals from that portfolio. The household may have other cash flows such as pensions or Social Security payments that cover essential needs. Consequently, variability in withdrawal amounts is not preferred but may be manageable if the disutility of this variability is properly accounted for. The retiree has settled upon a preferred allocation of these funds between risky and risk-free assets. The remaining question is, “What rule or policy will be used to decide how much of this portfolio to withdraw and spend each year?” When choosing among the infinite variety of such rules, having a consistent metric and a simple explanation of alternatives should prove useful.

A USEFUL METRIC: PART 1

If we wish to compare a collection of withdrawal strategies in a way that allows us to state that one strategy is better than another, we need to accomplish three things. First, we need to define a collection of scenarios over which our strategies will play out. Second, we need to specify the strategies in a consistent way. Third, we need to develop a way to describe the performance of each strategy relative to the others.

We evaluate the performance of a policy over a common set of 10,000 scenarios. Each scenario represents a $T = 30$ -year horizon. T can be higher or lower, but the methodology does not change. A scenario includes inflation-adjusted, or real returns, on a risky asset and a risk-free asset. We use the annual return of the S&P 500 as our proxy for risky assets, the annual returns on holding 30-day U.S. treasuries as our proxy

for risk-free assets, and the Consumer Price Index estimates as our proxy for inflation rates. All of this data is provided by Morningstar and covers the period 1926–2015. To be more precise, we randomly generate sets of 30 integers between 1926 and 2015. (We allow duplicates to exist within a set.) For each corresponding value we gather the return and inflation levels from those years. This matrix with 30 rows and 3 columns presents a single scenario. We generate 10,000 scenarios in this fashion and use this set in all the analysis that follows. We ignore transaction costs and tax effects to maintain the focus on our central ideas.

Given this construct, we can define a strategy using a pair of components: (1) s , a share of portfolio value, to be placed in the risky asset at the start of each year, and (2) a rule defining the proportion of portfolio value, p_t , withdrawn at the start of each year. Thus, s is a percentage of total value allocated to equities at the beginning of each year, and t indexes years within a scenario. An infinite variety of rules for defining p_t are possible including:

- select p such that the withdrawal amount equals 4 percent of the original portfolio balance,
- p is a fixed percentage of the portfolio balance at the start of the year,
- p is a fixed percentage applied to a moving average of account balances,
- p is selected in response to account performance,
- p is determined by RMD rules, etc.

To describe performance across an arbitrary set of strategies, we want to compare them all against a common benchmark that gives additional insight about how good a strategy is. One simple benchmark is found by calculating the utility generated from a strategy that withdraws a constant dollar amount each year to deplete the portfolio exactly at the end of the planning horizon. For each scenario, this value can be found, and it represents the maximum utility linked with level consumption over the planning horizon. The utility generated by any other strategy can be compared with this value. We label the ratio of these total utility values “efficiency.” When looking back at the end of the horizon, a policy with greater efficiency would be deemed better.

The term “utility” relates to the idea that withdrawal amounts are of no real value for their own sake but are converted to units of utility or happiness by the consumer. Numerically, this conversion is done via a preference function. The preference function also can account for the level of risk aversion of the decision-maker. When we consider the lengthy research stream on retirement planning, we find that when preference functions are formally considered, the vast majority of published works assume that the retiree exhibits constant relative risk aversion (CRRA). In simple terms, this is the assumption

that if an investor is indifferent between doing nothing and having a 50/50 chance at an increase of x percent versus a loss of y percent where $x > y$, then this evaluation is independent of the initial level of wealth. CRRA implies that if the utility increase from a 10-percent gain is equal and opposite to the utility decrease from a 5-percent loss, then this holds whether the investor starts with \$100 or \$1 million. Under this common assumption, it is easy to show that the only utility functions consistent with this story are transformations of $-\frac{1}{w^\gamma}$ where w represents a withdrawal amount and γ reflects the level of risk aversion (Bach 2010). Note that if this is the case then reaching a value of $w = 0$ is an outcome with utility of negative infinity.

Let us take a moment to consider one aspect of what this means. If retirees have CRRA, then implementing a policy that demands an annual withdrawal of a fixed amount is not rational. To see why this must be true, consider an initial balance of \$100 and annual withdrawals of \$4 over a 30-year horizon. Such strategies guarantee that there is a positive probability of exhausting the portfolio before the end of the planning horizon. Consequently, this approach assigns a positive probability to an event with a utility level of negative infinity, and the expected total utility of any such policy is also infinitely negative. For our setting, strategies that withdraw some fixed amount each year all have this characteristic and will not be modeled here. With this in mind, we assume that the decision-maker exhibits CRRA and write utility as,

$$U(w_t) = \frac{(w_t)^{1-\gamma}-1}{1-\gamma} \tag{1}$$

Given this preference function, we can easily translate a withdrawal amount into a utility level. If we have a stream of utility levels that stem from a stream of payments, we can find the average utility level and convert this to a level payment amount in real dollars that delivers the same total utility. This is referred to as the certainty equivalent withdrawal (CEW). To make these ideas more concrete, let us work through a small example.

ILLUSTRATIVE EXAMPLE, PART 1

Consider an investor with a three-year horizon starting with \$100 who withdraws 6 percent of the balance at the beginning of each year ($p = 6$ percent) and invests the rest in a portfolio

having a 50/50 equity/bond split ($s = 50$ percent). We seed the problem by drawing real investment results from three random years. In this example, these happen to be 1973, 1956, and 2003. We label the years as 1, 2, and 3, respectively, then add the starting balance, withdrawal amounts, total return, and ending balance to display the outcome as shown in table 1.

For the purpose of this example, let us set $\gamma = 2$. In this case, the utility levels of the three payments listed in table 1 become 0.833, 0.799, and 0.789. These values average 0.807. Thus, this stream of cash flows delivers a utility level of 0.807 units per year. The CEW is the annual cash flow offering the same level of utility. This is found by rearranging equation 1 to get,

$$CEW = (1 + U(1 - \gamma))^{1/(1 - \gamma)} \tag{2}$$

In equation 2, U is the average utility level of 0.807, and we find that $CEW = 5.18$.

To find efficiency we need to compare this value with the annual real dollar withdrawal amount that exhausts the portfolio after year 3. Prior works, including Suarez et al. (2015) and Blanchett et al. (2012), apply this idea and refer to this amount as the sustainable spending rate (SSR). Given the data in table 1, we see that the balance after year 3 will be,

$$B_3 = \{[(B_0 - w_1)(1 + r_1) - w_2](1 + r_2) - w_3\}(1 + r_3) \tag{3}$$

Here w_i and r_i are withdrawals and rates of return in year i . If we assume that $w_1 = w_2 = w_3$ and set $B_3 = \$0$, we can calculate this corresponding value of w , which is also the SSR. More generally, for a horizon of T years this value can be calculated as,

$$SSR = B_0 \times \frac{\prod_{t=1}^T (1 + r_t)}{\sum_{t=1}^t \prod_{j=1}^T (1 + r_j)} \tag{4}$$

where t and j are index years. Note that $\prod_{t=1}^T (1 + r_t)$ represents the product of T terms. For the scenario described in table 1, we get,

$$SSR = \$100 \frac{(1 - 0.115)(1 + 0.019)(1 + 0.121)}{(1 - 0.115)(1 + 0.019)(1 + 0.121) + (1 + 0.019)(1 + 0.121) + (1 + 0.121)} = \$30.87 \tag{5}$$

Table 1

EVOLUTION OF ACCOUNT BALANCE WITH REAL INVESTMENT RETURNS WITHDRAWING 6 PERCENT AT BEGINNING OF EACH YEAR FOR THREE YEARS

Year	Real S&P Return (%)	Real Bond Return (%)	Starting Balance	Real Withdrawal	Total Return	Ending Balance
1	-20.51	-2.54	100.00	6.00	-11.52	83.17
2	5.94	-3.76	90.53	5.43	1.09	79.03
3	26.06	-1.92	80.73	4.84	12.07	83.26

Given this CEW value and the SSR value we can state the efficiency of this strategy as the ratio of the two, or $\$5.18/\$30.87 = 0.1678$ or 16.78 percent.

In this instance the SSR values are less than one-third of the original principal even though the expected return on money invested at time 0 for three years would be roughly 1 percent. This occurs because the portfolio loses value during year 1 and the profitable returns in years 2 and 3 are insufficient to recover that loss. This highlights the fact that given some sequences of returns, alternate decision rules exist that would perform better than the benchmark.

A USEFUL METRIC: PART 2

Thus far we have compared the CEW with the SSR and define efficiency as $E = CEW/SSR$. In a world with perfect information and no variability, one can choose payments equal to the SSR and efficiency will be 100 percent. However, given risk aversion and uncertainty, this value will be less than 100 percent and greater values are inherently better. However, if we confine our strategy space to policies that leave a positive ending balance, this metric is incomplete because it needs to account for whatever utility level we assign to the ending balance.

We note that the level of risk aversion related to this terminal value may differ from that linked to annual withdrawals for the purpose of consumption, and it seems reasonable to assume that this level will be lower. With this in mind, we write the utility of this amount using the same functional form as in equation 1 to get $U(B_T) = [(B_T)^{1-\alpha} - 1]/(1-\alpha)$, noting that α is less than or equal to γ . We now expand our notion of efficiency and define a more complete metric that we label relative efficiency or RE as,

$$RE = \frac{CEW + U(B_T)}{SSR} \tag{6}$$

Note that when we associate some utility with the terminal value, we acknowledge that the total utility generated exceeds that which is brought about by spending over the planning horizon. This is reasonable because leaving some of the funds as a bequest was an explicit desire and some utility should be linked to fulfilling that desire. As a consequence, it is possible for this metric to exceed 100 percent.

After we find the relative efficiency of a policy over each scenario, we can consider how this plays out over the entire set of scenarios. This distribution of RE values can be analyzed in many ways. For now, we choose to focus on two simple metrics: the expected value of relative efficiency and the likelihood that a strategy exceeds the benchmark. We label these as mean relative efficiency or MRE, and a likelihood of exceeding the benchmark or LEB. These are calculated as,

$$MRE = \frac{1}{N} \sum_{n=1}^N RE_n \tag{7}$$

$$LEB = \frac{\# \text{ of times } RE > 1}{N}$$

Note that n is being used as a counter over our $N = 10,000$ scenarios.

ILLUSTRATIVE EXAMPLE, PART 2

Returning to the example laid out in table 1, recall that we found an SSR value of $\$30.87$. For the policy under consideration where $(p, s) = (6 \text{ percent}, 50 \text{ percent})$, we calculated an average annual utility level of 0.807 and an ending balance of $\$83.26$. To complete the analysis, we calculate a utility level for the ending balance using equation 1 to get $U(B_3) = 0.988$. We insert these values into equation 6 to get $RE = (5.18 + 0.988)/30.87 = 19.98 \text{ percent}$. Note that the CEW of $\$5.18$ is below the average payment of $(\$6.00 + \$5.43 + \$4.84)/3 = \5.42 . This must be the case because the investor dislikes variability in payout amounts.

In addition, we note that the utility linked to the ending balance is small relative to that balance. In effect, we are dramatically discounting the utility derived from the terminal value. If a client wishes to put more weight on this value, it can be multiplied by a constant equal to more than 1. However, our intent is to account for this value without having it dominate our results, so we set this constant at 1 in the discussion that follows.

SIMPLE WITHDRAWAL POLICIES

To build intuition, we consider a series of increasingly complex policies to determine withdrawal amounts from an evolving portfolio. The simplest policy that we focus on involves liquidating a fixed percentage of asset value annually. Table 2 shows MRE values when $s = 1$, no utility is assigned to B_{30} ,

Table 2 MRE LEVELS WITH $4\% \leq p \leq 10\%$ AND $2 \leq \gamma \leq 6$

$\gamma \downarrow, p \rightarrow$	4%	5%	6%	7%	8%	9%	10%
2	74.79	81.48	84.51	84.49	82.00	77.62	71.88
3	68.10	74.90	78.10	78.16	75.60	71.00	64.94
4	63.33	70.04	73.30	73.44	70.92	66.33	60.25
5	59.81	66.39	69.64	69.83	67.40	62.88	56.90
6	57.12	63.56	66.78	67.01	64.66	60.25	54.38

Table 3

MRE AND LEB LEVELS WHEN $p = 6.5\%$, $2 \leq \alpha \leq 6$ AND $2 \leq \gamma \leq 6$

$\gamma \downarrow, \alpha \rightarrow$	2	3	4	5	6
2	101.70 54.14	93.46 29.97	90.34 23.99	88.20 21.46	85.16 20.02
3	95.34 41.83	87.10 13.79	83.99 5.69	81.84 0	78.80 0
4	90.58 19.25	82.34 0.10	79.26 0	77.08 0	74.04 0
5	86.95 7.36	78.71 0.05	75.59 0	73.45 0	70.41 0
6	84.10 4.03	75.86 0.04	72.74 0	70.60 0	67.56 0

γ ranges from 2 to 6, and p ranges from 4 percent to 10 percent. We see an MRE value of 73.30 when $\gamma = 4$ and $p = 6$ percent. This means that, on average, the total utility offered by this strategy is roughly 73 percent as high as could have been obtained using a level payout amount under perfect information. We sometimes refer to such approaches as “endowment strategies” because they are similar to the way that many endowments are managed. The shaded columns in table 2 highlight the fact that MRE values are maximized between 6 percent and 7 percent for all γ levels considered.

Recall that this is the same strategy used in the earlier example, but the efficiency level is dramatically higher. For a fixed value of p , lengthening the planning horizon increases efficiency because the accumulated annual withdrawals increase as a proportion of the initial portfolio value with the planning horizon. We also see that increasing the level of risk aversion reduces efficiency noticeably but does not change the fact that the most attractive withdrawal rate is between 6 and 7 percent. As γ values rise, MRE levels decline because being more sensitive to variability is synonymous with saying that variability reduces efficiency. Setting p to 6 percent we see MRE drop from 84.5 percent to 66.78 percent as γ rises from 2 to 6. To put this in perspective, setting γ to 2 means that increasing the amount of a withdrawal by 50 percent, say from \$4 to \$6, increases

utility by roughly 11 percent. Reducing the amount by 50 percent drops utility by 33 percent. Thus a 50-percent loss is three times as painful as the joy of a 50-percent gain. Empirical evidence suggests that γ values around 4 are consistent with the behavior of many investors (Azar 2006). If this is the case, then the pain of a 50-percent drop is roughly 10 times as great as the joy of a 50-percent gain.

Table 3 reports MRE and LEB values when $p = 6.5$ percent, the utility of B_{30} is included, and both α and γ values range from 2 to 6. These MRE levels are uniformly greater than corresponding values in table 2 because they reflect the added utility of a positive ending balance. Note that when $\alpha = \gamma = 2$, we have $MRE > 1$ and $LEB > 50$ percent. In other words, this policy is likely to outperform the benchmark if this set of parameters is in place. It also turns out that as long as both α and γ values are between 2 and 6, the most attractive value of p (labeled p^*) is always between 6.0 percent and 7.0 percent. Thus, the selection of p is not very sensitive to the level of risk aversion. This is important because estimation of these values is not as easy as it may seem and attitudes toward risk may evolve over time based on a host of factors including portfolio returns, health, and major life events. For convenience let us label the scenario with $\alpha = 2$, $\gamma = 4$, withdrawals defined by a single value of $p_t = p$ as our Base Case.

Figure 1

RE VALUES UNDER BASE CASE WHEN $s = 100\%$

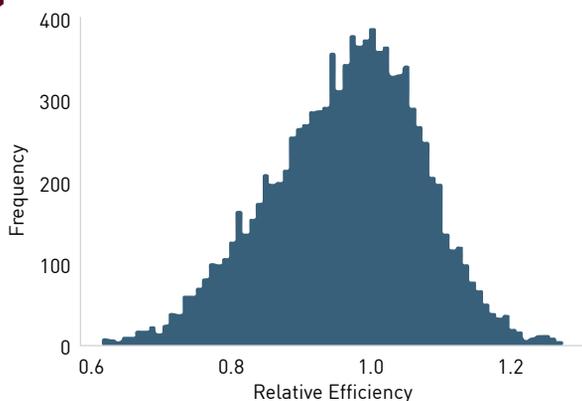


Figure 1 shows the histogram of RE values for the Base Case with $p = 6.5$ percent and $s = 100$ percent.

Figure 2 shows MRE values as functions of p for portfolios with $s = 20$ percent, 40 percent, 60 percent, 80 percent, and 100 percent. Here we see that the level of p that maximizes MRE does change slightly with the value of s . Considering s values of 60 percent, 80 percent, and 100 percent, we see peak MRE values of 90.01, 91.09, and 90.58 that occur at $p = 5.5$ percent, 6.0 percent, and 6.5 percent, respectively. For this set of portfolios efficiency is maximized with $s = 80$ percent, and $p = 6.0$ percent. This can be compared with the classic 60/40 split, which performs very close to this optimized portfolio and policy.

EFFICIENCY AND THE SEQUENCE OF RETURNS

When withdrawals are made over a sequence of periods, the order of returns matters. Sequence of returns risk relates in part to the fact that if an investor has several bad years at the start of the planning horizon, it has a disproportionate impact on the total utility generated. This issue was illustrated in the example presented in table 1. To provide some insight into this phenomenon, figure 3 displays MRE values for the Base Case with an additional factor included: The returns for the first three years of each scenario are all artificially reduced by 10 percent.

For this setting the maximum MRE value shown occurs when $p = 6$ percent and $s = 80$ percent just as was the case in figure 2. However, the best MRE value is noticeably higher in figure 3 than the maximum shown in figure 2, 95.07 percent versus 91.09 percent. This is not to say that net utility is higher. Clearly it is reduced when returns are lower. The key observation is that the policy considered here gets closer to the benchmark when the first few years of returns are reduced and the best (p, s) pair is the same. This happens because policies built around proportional withdrawals reduce the drawdown when the portfolio underperforms. Having the ability to accept lower payments given disappointing returns is akin to purchasing a hedge against downside risks.

Sequence of returns risk relates in part to the fact that if an investor has several bad years at the start of the planning horizon, it has a disproportionate impact on the total utility generated.

EFFICIENCY AND THE PLANNING HORIZON

When assets are pooled to serve a large population of retirees, it is inevitable that some retirees will withdraw funds over a shorter horizon and others will make more withdrawals than expected. Pooling assets in a pension fund effectively means that those who receive fewer than the expected number of payments effectively subsidize those who receive more. As a result, a central planner can manage the portfolio as though it is serving a group of people who will all have roughly an average life-span. When an individual manages a personal account, this risk-pooling feature is not available. Consequently, retirees are forced to consider longer planning horizons.

Figure 4 shows MRE levels for a range of p values under the Base Case, with s set to 80 percent and the planning horizon set to 20, 30, and 40 years. The shapes of these curves suggest that as the horizon becomes longer, the peak efficiency of these

Figure 2

MRE VALUES UNDER THE BASE CASE

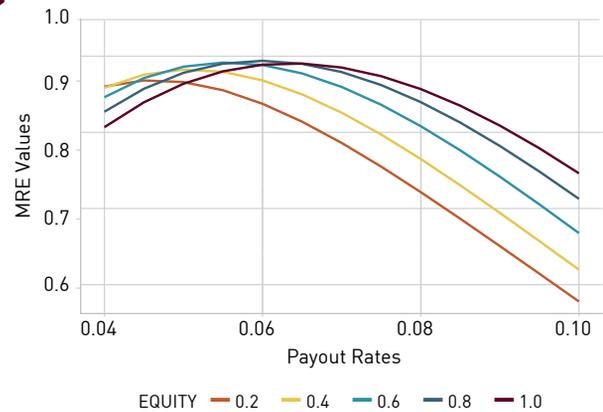


Figure 3

MRE VALUES FOR BASE CASE WHEN $R_1, R_2,$ AND R_3 ARE EACH REDUCED BY 10 PERCENT

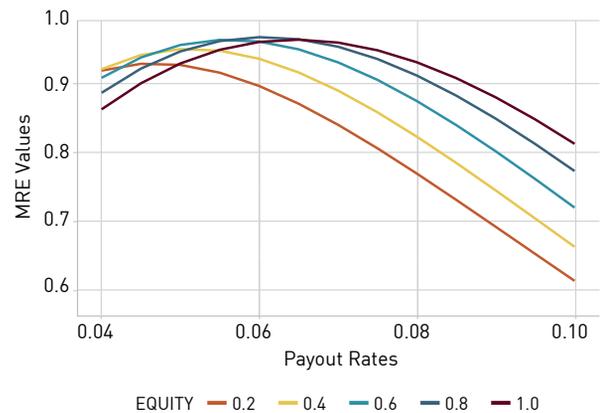
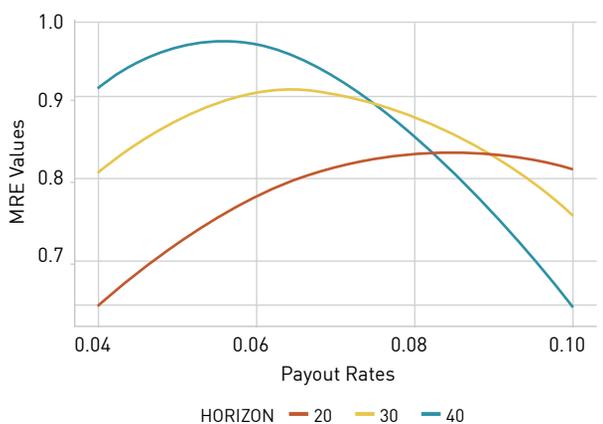


Figure 4

MRE VALUES UNDER THE BASE CASE WITH $T = 20, 30,$ AND 40



policies rises and the payout rate that maximizes efficiency falls. This is consistent with the intuition that planning for a longer horizon should be matched to reduced withdrawal amounts to hedge against future difficulties.

Table 4

MRE AND LEB LEVELS USING RMD RULES WHEN $s = 80\%$, $2 \leq \alpha \leq 6$ AND $2 \leq \gamma \leq 6$

$\gamma \downarrow, \alpha \rightarrow$	2	3	4	5	6
2	111.08 99.52	102.95 80.94	100.03 53.53	98.55 36.73	97.65 28.18
3	102.95 64.46	94.83 18.70	94.90 2.52	90.42 0.27	89.53 0.03
4	96.95 42.54	88.83 8.07	85.90 0.59	84.42 0.03	83.53 0
5	92.49 31.55	84.36 4.23	81.44 0.20	79.96 0	79.06 0
6	89.08 25.08	80.96 2.22	78.03 0.08	76.55 0	75.66 0

Table 5

VARIATIONS OF RMD RULES FOR BASE CASE

Rule	$s = 60\%$	$s = 80\%$	$s = 100\%$
RMD	101.85 61.98	96.95 42.54	92.39 31.16
RMD-Plus	103.89 72.40	99.67 54.34	95.25 38.33
RMD-Minus	98.19 45.70	92.93 30.95	88.42 23.44
Inverse RMD	46.12 0.00	49.27 0.06	52.41 1.02

Equivalently, these same curves also support the contrapositive position, meaning that when one has a shorter planning horizon the payout rate that maximizes efficiency increases. This strongly suggests that policies that have payout rates increase over time will outperform strategies in which a payout rate is selected and never changes.

DYNAMIC WITHDRAWAL RATES

Rules using a fixed withdrawal rate are simple to explain and are quite efficient even when early returns are reduced. However, such policies have two notable shortcomings. First, they leave most of the money on the table. If retirees never remove more than 6-7 percent of portfolio value, then the bulk of the funds available is never used by the investor. Second, intuition strongly suggests that policies that account for performance or a shortening planning horizon should outperform simple static rules.

To explore these issues, we compare several additional policies. Let us refer to the policy that sets p to a single value as the Single Value Rule. Define Rule 1 as a policy in which the payout is a percentage of a moving average of the ending balances over the three most recent years. This naturally smooths the outlays, and the concavity of the utility function suggests that this may increase efficiency. Rule 2 is a policy in which p_d is the share withdrawn if portfolio value dropped over the past year and $p_r = p_d - 1$ is the share withdrawn if that value rises. Having p_r to be strictly less than p_d serves to reduce the variability of payouts and may offer an advantage in that withdrawing a smaller percentage when values rise retains more of the

gain for later use. Table 4 displays MRE values using all three rules under the Base Case when $s = 80$ percent.

The performance of each of these rules is rather similar and the most efficient value of p is roughly 6 percent in all cases. Using the Single Value Rule, the maximum MRE value is 91.06 percent. The maximum MRE value for Rule 1 is 90.61 percent. Thus, we see that Rule 1 underperforms the Single Value Rule. An exploration of a variety of alternate smoothing rules offers similar results. It appears that because these approaches do not allow the amount withdrawn to adjust as quickly as the Single Value Rule, some efficiency is lost. On the other hand, under Rule 2 the best MRE value rises to 91.17 percent. This approach removes the lag between a drop in value and the reaction and allows for an additional level of adjustment to preserve gains for future use.

In our search for policies providing a clearly better performance, we turn to several variants of approaches based upon RMD rules. Table 4 displays results for a simple rule built around RMD guidelines. Specifically, this rule applies the RMD framework in the following way. The withdrawal rate in year 1 is the same as dictated by the RMD rules that kick in at age 72. The rate for year 2 is what the RMD rule dictates for age 73, and so on. This rule is a bit artificial in that many retirees begin withdrawals long before this age. However, this does present a useful comparison because so many retirees do make withdrawals in accordance with RMD rules by choice or otherwise. Under this arrangement, the percentage of the portfolio that is liquidated in each year rises steadily from $1/27.4 = 3.6$ percent in year 1 to $1/5.9 = 16.9$ percent in year 30. Each cell of table 4 shows MRE values on the top line, with the corresponding LEB value below it for the same range of risk aversion parameters considered in table 3. For example, for the Base Case in table 4 the MRE value is 96.95 percent compared to 90.58 percent from table 3. In addition, this policy outperforms the benchmark 42.54 percent of the time compared to 19.25 percent of the time for the policy depicted in table 3. In fact, comparing table 4 to table 3 shows that this RMD-based rule outperforms our Single Value Rule in every scenario tested and on average the performance gap is 7.9 percent (89.60 percent versus

81.70 percent) for the MRE and 10.31 percent (19.98 percent versus 9.67 percent) for LEB.

Although this RMD rule clearly outperforms the simpler policies, it is worthwhile to consider results for several other variants as well. Table 5 shows MRE and LEB values for the basic RMD rule on the top row. The row labeled “RMD-Plus” refers to a policy in which the corresponding withdrawal rates are all multiplied by 110 percent. This raises the withdrawal in year 1 to roughly 4 percent. The row labeled “RMD-Minus” multiplies the withdrawal rates by 90 percent. Finally, “Inverse RMD” refers to a policy that inverts the withdrawal rates. In other words, it starts with a 16.9-percent rate in year 1 and walks down to 3.6 percent in year 30. MRE values when using RMD-Plus are roughly 2 percent greater than those found using the simple RMD rule, which in turn are roughly 2 percent above those for the RMD-Minus rule. This declining withdrawal rate is consistent with spending that declines over time which is commonly seen in practice (Roy and Hahn 2021).

In short, our results show that rules based on RMD requirements are surprisingly efficient and there is some evidence that being even more aggressive increases efficiency. This does not mean that such rules are perfect or cannot be improved with a simple adjustment. However, it does demonstrate an upside to such policies that should not be ignored.

DISCUSSION AND CONCLUSIONS

When helping decision-makers with hard problems, it is often useful to employ simple metrics to summarize a large body of data and its analysis. Avoiding information overload has clear advantages. The notion of efficiency is understood as a quick and dirty way to compare units in some set. Our application of this idea boils the results of a comparative analysis into a single value. This metric subsumes the presentation of a great deal of data and is constructed with information about the decision-maker’s portfolio, wishes, and attitudes toward risk.

We note that this idea is not new to the field. The focus on efficiency within a constant relative risk aversion framework has been used by others including Blanchett et al. (2012), Williams and Finke (2011), and Delorme (2015a, b). Several conclusions are common across these works. First, some dynamic approaches will outperform static policies. This is not surprising given the incorporation of risky assets. Second, results are not very sensitive to the specific risk aversion level in place (Delorme 2015a, Blanchett et al. 2012). This is important because these parameters can be hard to estimate and may not be constant across the problem horizon. Third, the presence of alternate income streams can make portfolios with higher equity portions and higher withdrawals more efficient. (Finke et al. 2012; Milevsky and Huang 2011; Williams and Finke 2011). At the same time, this work adds consideration of the utility that stems from the remaining value at the end of the

planning horizon. This adds an important dimension to the discussion and forces us to think about how often a policy actually exceeds the benchmark.

On the other hand, whenever a single value is used to summarize a more complex analysis this comes at some cost, and we must be concerned about the distortions introduced. A key question is whether an approach that seeks to simplify policy comparison is likely to lead to behaviors that are opposed to the decision-maker’s best interest. We would be concerned if a simplification masks the risks involved or takes attention away from important issues. The results here assume that the decision-maker’s level of risk aversion is known. Fortunately, policy choice is not very sensitive to the level of risk aversion. This reduces the cost of not being able to find this level precisely. In addition, this work focused on investors who are comfortable with variable withdrawals from the accounts considered. This is key because it mitigates sequence of returns risk. If that is not the case, then alternate approaches should be considered. For an exhaustive discussion of withdrawal rules see Pfau (2017).

When helping decision-makers with hard problems, it is often useful to employ simple metrics to summarize a large body of data and its analysis.

Finally, longevity risk is important, especially if a single withdrawal rate is sought. Fortunately, this was accounted for when RMD calculations were developed. For these reasons, it seems fair to say that the approach presented here should not increase the likelihood of risky behavior that is not properly accounted for.

With this being said, there is a major issue that we have left largely undiscussed to this point. Most retirees will enter any conversation about planning for withdrawals having heard of simple approaches such as the famous 4-percent rule, and some will argue that such policies get rid of this pesky variability of payments (Bengen 1994). However, it is irresponsible to ignore the fact that under the 4-percent rule, a 4-percent withdrawal rate only applies to the first year. If the portfolio declines in value, the 4-percent rule can result in withdrawing much more than 4 percent of portfolio value in subsequent periods. In the worst case, it results in withdrawal of 100 percent of that value, leaving nothing behind.

The policies considered here do not eliminate risk. But neither does any other approach if a risky asset is being held over time. When comparing the approach suggested here with older approaches such as the 4-percent rule, the key observation is

that the nature of the risk differs. Ignoring the variable nature of withdrawal levels on a percentage basis by focusing only on certainty of withdrawal amounts on an arithmetic basis does not mean that rules based on a fixed dollar amount are less risky. In addition, the risk in the approaches we lay out here is being formally included in the analysis, and this is not possible with some other approaches without artificially assigning some finite utility level to an empty account. Thus, we can say that instead of increasing uncertainty, the approach that we suggest here exchanges a hidden uncertainty that is unaccounted for with a revealed uncertainty that is fully accounted for.

Withdrawal rules based on withdrawal percentages are easy to explain and surprisingly efficient. On the other hand, we also see that as the horizon gets shorter, the withdrawal percentage that maximizes efficiency rises. Consequently, having a policy with a withdrawal rate that rises over time performs better. With this in mind, we explored a family of policies built around rules used to dictate required minimum distributions. It turns out that these rules are much more efficient than even the best single withdrawal rate policy that we could find and this is true for all of the settings that we studied. Although we would never claim that this is the best policy, the central message remains that endowment-style policies are worthy of more serious consideration than they have received to date and building such approaches around widely known rules such as RMDs is at least a sound strategy. In addition, it is useful to explore the utility generated by RMD rules because recent survey results verify that the majority of retirees with retirement accounts actually make account withdrawals based on the RMD rules (Roy and Hahn 2021). Because many retirees choose to (or are effectively forced to) work with RMD rules anyway, providing a narrative that the efficiency of such strategies is surprisingly high may help manage the anger that RMD rules generate.

In closing, many investors have a natural interest in approaches to managing a portfolio of assets used to support retirement years in a way that avoids even the possibility of exhausting

these funds. Strategies based on fixed or growing withdrawal rates are one family of approaches that facilitates this in a way that is easy to explain and implement. Consequently, investors and advisors should benefit from our presentation concerning the utility of such approaches as guideposts in the efforts to maximize a retiree's standard of living. ●

Chester Chambers, PhD, is Professor of Practice and Academic Program Director for Business Analytics and Risk Management Programs at the Johns Hopkins Carey Business School. Contact him at cchamber@jhu.edu.

REFERENCES

- Azar, S. M. 2006. Measuring Relative Risk Aversion. *Applied Financial Economics Letters* 2, no. 5: 341–345.
- Back, K. E. 2010. *Asset Pricing and Portfolio Choice Theory*. New York: Oxford University Press.
- Bengen, W. P. 1994. Determining Withdrawal Rates Using Historical Data. *Journal of Financial Planning* 7, no. 4 (October): 171–180.
- Blanchett, D., M. Kowara, and P. Chen. 2012. Optimal Withdrawal Strategy for Retirement-Income Portfolios. *Retirement Management Journal* 2, no. 3: 7–20.
- Delorme, L. 2015a. Confirming the Value of Rising Equity Glide Paths: Evidence from a Utility Model. *Journal of Financial Planning* 28, no. 5: 46–52.
- . 2015b. A Blueprint for Retirement Spending. *Journal of Financial Planning* 28, no. 9: 40–50.
- Finke, M., W. D. Pfau, and D. Williams. 2012. Spending Flexibility and Safe Withdrawal Rates. *Journal of Financial Planning* 25, no. 3: 44–51.
- Milevsky, M. A., and H. Huang. 2011. Spending Retirement on Planet Vulcan: The Impact of Longevity Risk Aversion on Optimal Withdrawal Rates. *Financial Analysts Journal* 67, no. 2: 45–60.
- Pfau, W. 2017. *How Much Can I Spend in Retirement? A Guide to Investment-Based Retirement Income Strategies*. McLean, Virginia: Retirement Researcher Media.
- Ritholtz, B. 2017. Tackling the Nastiest, Hardest Problem in Finance. *Bloomberg View* (June 5). <https://www.bloomberg.com/opinion/articles/2017-06-05/tackling-the-nastiest-hardest-problem-in-finance>.
- Roy, K., and K. Hahn. 2021. Mystery No More: Portfolio Allocation, Income and Spending in Retirement. JP Morgan (August 19). <https://am.jpmorgan.com/us/en/asset-management/adv/insights/retirement-insights/retirement-portfolio-allocation/>.
- Suarez, E. D., A. Suarez, and D. T. Walz. 2015. The Perfect Withdrawal Amount: A Methodology for Creating Retirement Account Distribution Strategies (August 3). <https://ssrn.com/abstract=2551370>.
- Williams, D., and M. Finke. 2011. Determining Optimal Withdrawal Rates: An Economic Approach. *Retirement Management Journal* 1, no. 2: 35–46.



INVESTMENTS & WEALTH INSTITUTE®

5619 DTC Parkway, Suite 500
Greenwood Village, CO 80111
Phone: +1 303-770-3377
Fax: +1 303-770-1812
www.investmentsandwealth.org

© 2021 Investments & Wealth Institute®. Reprinted with permission. All rights reserved.

INVESTMENTS & WEALTH INSTITUTE® is a registered mark of Investment Management Consultants Association Inc. doing business as Investments & Wealth Institute. CIMA®, CERTIFIED INVESTMENT MANAGEMENT ANALYST®, CIMC®, CPWA®, CERTIFIED PRIVATE WEALTH ADVISOR®, RMA®, and RETIREMENT MANAGEMENT ADVISOR® are registered certification marks of Investment Management Consultants Association Inc. doing business as Investments & Wealth Institute.