

Predicting the Success of a Retirement Plan Based on Early Performance of Investments

CS229 Autumn 2010 Final Project

Darrell Cain, AJ Minich

Abstract

Using historical data on the stock market, it is possible to predict the historical success rates of given retirement plans. The fundamental problem with retirement planning is the inability to collect data on the performance of the investments in the future. Thus a retiree often does not know whether or not his plan will succeed or fail until they are well into the plan itself. In this paper, we address the effectiveness of assessing a retirement plan based on the first few years of market performance.

Introduction

Many wage-earners face great uncertainty upon entering retirement: even with enormous savings, their own futures and the future of investment markets is impossible to predict. Given the wide number of variables affecting their portfolio's performance, retirement planning is often a long-term bet based on very little information. A retiree's worst financial nightmare, of course, is running out of savings before his end of life, so the consequences are dire if he places the wrong.

To capture the predicament facing retirees, we develop an equation called the initial withdrawal rate (IWR) equation which gives the percentage amount of the initial savings that can be spent over each year (taking inflation into account). This equation predicts the annual yearly buying power of a retiree over the course of their retirement:

$$IWR_r = \frac{\prod_{i=1}^{r-1}(1+g_i)}{\sum_{i=1}^{r-2}[\prod_{j=1}^i(1+in_j)\prod_{j=i+1}^{r-2}(1+g_j)] + \prod_{i=1}^{r-1}(1+in_j)}$$

where g_i is the growth of the assets in period i , in_i the inflation in period i , r the number of retirement years, F_r the amount of money left over at the end of retirement, and I_0 the initial savings. (See **Appendix: Calculating SWdR** for a complete derivation.)

Ideally, the retiree would like to select a rate that results in the funds being nearly completely exhausted at the end of his or her life. (For the purposes of this paper, we will assume that the retirement span r is exactly 30 years.) By setting the final amount F_r to zero, we can calculate a special IWR called the **Safe Withdrawal Rate**, or SWdR. The SWdR represents the maximum amount a retiree can spend on a yearly basis without borrowing – that is, the optimal balance between retirement lifestyle and financial security.

This concept raises perhaps the most important question in retirement: how does one decide his or her own SWdR? After all, the equation for IWR (and thus SWdR) depends on growth and inflation rates throughout the retirement, which the retiree obviously does not know at the beginning of retirement. There is also the issue that, historically, the variance can be quite large: the SWdR reaches as high as 15% in 'good' times (such as 1950-1980) and as low as 4% in depressed economies. Selecting a low SWdR hedges risk, but has a significantly negative impact on the 'golden years' that the retiree has worked so hard to earn.

The current approach was pioneered in 1994 by Bengen¹, and involves calculating the historical SWdRs for an asset and then choosing the minimum of those values. The approach relies on the idea that the market will perform no *worse* than it has at some point in the past. For years, financial planners have used Bengen's method to attain the 'Rule of 4%', which (as mentioned above) is the absolute minimum of all SWdRs in the period over which financial data is available (1926-2010). Although this approach promises 100% certainty of a successful retirement, it is rather conservative, requiring the retiree to live as if their retirement will span the worst economy in the last century. Thus in the vast majority of cases, the Rule of 4% leaves the retiree with significant savings at the end of retirement (experimentally, this can be as much as 14

¹ William P. Bengen, [Determining Withdrawal Rates Using Historical Data](#), Journal of Financial Planning, October 1994, pp. 14–24.

times the initial portfolio value). As an added difficulty, the retiree will often adjust the yearly withdrawal rate based on previous years' rates, which increases the retirement-plan SWdR but adds nonlinear terms to the IWR equation. Bengen's method cannot account for these sophisticated year-to-year relationships, generating suboptimal numbers even when the retiree performs the most basic withdrawal adjustments.

Though the retiree can increase his withdrawal rate above Bengen's conservative estimate during the retirement, he will need to live frugally in the early years of retirement and only increase his withdrawal rate near the end of the plan. This situation, too, is suboptimal: the retiree would prefer to live above his means during the first few years of retirement and make adjustments after a certain number of years.

We aim to develop and test an algorithm whereby a retiree could predict his retirement plan's full-length SWdR after only five years of retirement. If it is possible to predict the SWdR for a given retirement before the retirement is complete, then the retiree can adjust his income within a few years of entering retirement. For the purposes of this paper, we will consider the algorithm successful if, with only five years of retirement portfolio financial history, it predicts the 30-year SWdR within 20% of the true value with 90% confidence.

Creating a Dataset

We start by building a dataset on which to test our algorithms. We define the set of input features to be the following parameters:

- Growth data g of each asset for each year of the retirement
- Inflation data inf for each year of the retirement
- Average return r of the entire portfolio for the 30-year period
- Standard deviation σ of the portfolio for the 30-year period

We add the last two features primarily to differentiate between different portfolios which may have the same types of assets, but differing amounts of each.

Thus for an n -year retirement period beginning in year i , we define our vector of input features as

$$\hat{x}_i = [1 \quad r_i \quad \sigma_i \quad g_i \quad inf_i \quad \cdots \quad g_{i+n-1} \quad inf_{i+n-1}]$$

The output results are the actual Safe Withdrawal Rates for the 30-year periods corresponding to the data sets. Thus we let $y_i = \text{SWdR}$ for 30 year period in year i .

Given all input features x_i , one can calculate the exact SWdR y_i using the equations discussed above. However, we focus on attaining sufficiently accurate SWdR when we limit the features down to only the first $m < n$ years of growth and inflation data. Thus we define

$$x_{mi} = [1 \quad r_i \quad \sigma_i g_i \quad inf_i \cdots g_{i+m-1} \quad inf_{i+m-1}]$$

Thus $x_{5,6}$ would be a set of input features containing 5 years of growth/inflation data, and based on a retirement portfolio beginning in year 6 of our historical data (which is 1932, given that our data begins in 1926).

Modeling with Linear Regression

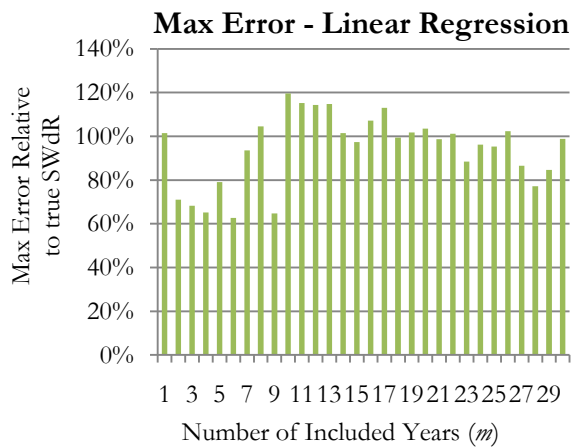

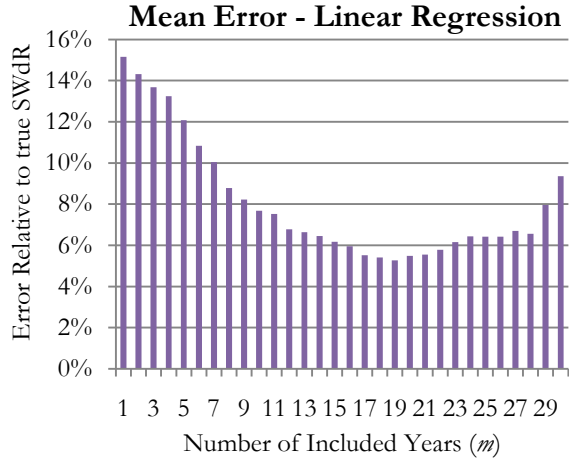
First, we assume that the data follows a fairly linear relationship, allowing us to use linear regression to generate a model of the data. We select a value $m \leq n$, define X to be a matrix containing successive x_{mi} on each row, and Y as a vector containing the y_i values for those corresponding years. In creating our linear regression model, we are looking for a hypothesis vector θ that satisfies $Y = X\theta$ (we have included a y intercept in our definition of x_{mi}). Thus as long as the training set contains more than $3 + 2m$ examples, we can solve our model using least-squares approximation:

$$\theta = (X^T X)^{-1} X^T Y$$

For each value of $m = \{1, 2, \dots, 30\}$, we used a training set composed of 100 different portfolios beginning in different years to generate our model, and then ran our prediction on a test set containing $k = 10,000$ such different portfolios. We define the error ($\epsilon^{(i)}$), mean error ($\bar{\epsilon}$) and the failure rate (ϵ_r) as follows:

$$\epsilon^{(i)} = \frac{\|y^{(i)} - y_{ls}^{(i)}\|}{y^{(i)}} \quad \bar{\epsilon} = \frac{1}{k} \sum_{i=1}^k \epsilon^{(i)}$$

$$\epsilon_r = \frac{\sum_{i=1}^k \mathbf{1}[\epsilon_i > .2]}{k}$$



The figures above show the average error, the failure rate and the maximum error across m . We can see that the mean error reflects generally good performance, with relative accuracy dropping to within 10% of the true value after only $m = 7$ years. The failure rate also reveals reasonable performance, with fewer than 10% of instances resulting in poor

performance when more than seven years of data are available.

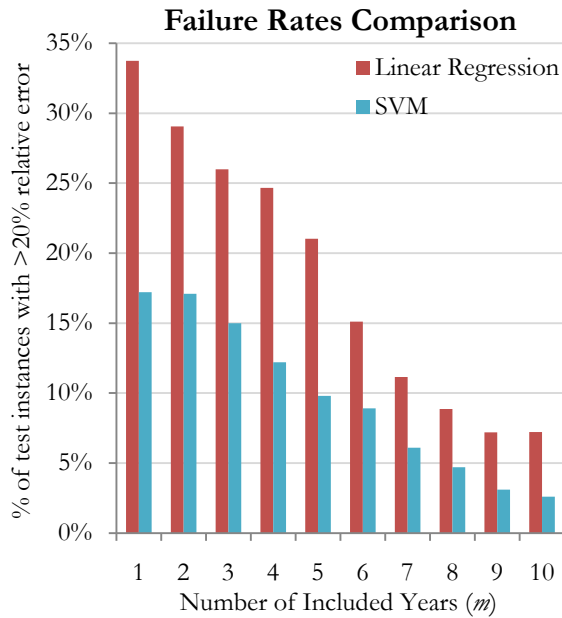


However, when we investigate the size of the maximum errors affecting both the mean error and the failure rate, we see that the model can generate truly useless results. In 10% of all instances, linear regression may predict a SWdR twice as much as the actual value, or it may predict that the retiree’s SWdR is *negative* – he must continue adding money to fund his ‘retirement’. While useful for giving us a baseline for performance, linear regression fails to provide sufficiently precise answers.

Modeling with a Support Vector Machine

Linear regression fails in part because it expects the values to lie along a hyperplane, despite the inherent nonlinearity of our problem. A more sophisticated fitting algorithm, such as a support vector machine (SVM), could model these relationships and predict SWdR values in a high-dimensional feature space. Additionally, an SVM fits naturally with our intuition that only certain growth and inflation rates in our input feature set – particularly the high-gain and high-loss years – will have a discernible impact on the actual SWdR value, and will thus become our support vectors.

Since an SVM mainly works as a binary classifier, we will modify the problem statement as follows: a given withdrawal rate can be either above or below the SWdR, indicating the potential success or failure as a retirement plan. In this case, the SWdR serves as a boundary value between the set of all successful and failure withdrawal rates. We choose a set of l evenly spaced withdrawal rates above and below the SWdR, and include these as input features along with the corresponding growth and inflation rates. For each withdrawal rate below the SWdR, the classifier is +1 (indicating a successful withdrawal rate), and for each rate above the SWdR, it is -1 (indicating failure). After training the model and running a prediction, we identify the original SWdR by feeding in several withdrawal rates along with the growth and inflation conditions and identifying where the boundary lies.

It is important to note that, while this method results in a better model, its error depends not only on the classification success but also the spacing λ between the withdrawal rates used to create the new



feature set. Even when the SVM properly classifies the boundary value, the resulting SWdR prediction can be as much as $\lambda/2$. Thus we will define the SVM's error not as the ratio of misclassified training examples, but as the error between the predicted boundary value and the actual SWdR. To minimize this value, we will declare a very small step size λ (typically around .0001, or .01%).

The results for the described SVM² are shown above. The method achieves excellent mean error, falling below 6% after 5 years. The failure rate results, compared above to those of linear regression, also

satisfy our requirements: with $m \geq 5$, the cases of failure drops below 10%. Thus with the SVM, we have successfully predicted SWdR within the desired tolerance given only the first five years of retirement portfolio performance.

Similarly, the maximum error with the SVM is well below that of linear regression, especially for $m \geq 5$ when it drops below 40% worst-case. This error bound represents a significant improvement over linear regression, and proves that we may be confident in the results generated using the SVM.

Conclusion

Although linear regression failed to provide a successful prediction algorithm, a standard SVM with a linear kernel met our criteria. With the SVM, we can predict SWdR to within 20% in 90% of test cases, and to within 40% in all of the test cases. Although we imagine that a retiree would want better accuracy for the purposes of retirement planning, we point out that the retiree must take on *some* risk in order to maximize the annual withdrawal rate. Using our proposed method and taking on minimal risk, the retiree may (in some cases) withdraw 50% more from his portfolio annually than when using Bengen's approach to calculating the SWdR. With those additional funds, the retiree can enjoy more of the retirement funds he has worked so hard during his life to earn.

² R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A Library for Large Linear Classification, *Journal of Machine Learning Research* 9 (2008), 1871-1874. Software available at <http://www.csie.ntu.edu.tw/~cjlin/liblinear>.

Appendix: Calculating SWdR

Originally written by Darrell Cain for Cain Watters and Associates. Utilizes historical financial data from IBB to calculate appropriate SWdR, given an initial portfolio and the years across which the retirement lasts.

At the start of retirement an individual has an initial amount of savings. At any given point in time the individual will withdraw a specific amount of money from that savings.

For all analysis done the point of withdrawal is chosen to be the end of the year. From this the *current withdrawal rate (CWR)* is defined.

$$CWR_i = \frac{W_i}{I_i}$$

The CWR for year i is given as the withdrawal amount W in year i divided by the amount of money I the retiree has in year i .

The goal of retirement is to predict the buying power needed by the retiree for each year and make sure that amount is available that year.

Each year the remaining amount of money grows by growth rate g . This gives the following equation for the amount of money left at year i . Note that withdrawals are made at the end of the year, this will be addressed later.

$$F_i = ((I_0(1 + g_1) - W_1)(1 + g_2) - W_2) \dots W_{i-1}$$

While this equation is fairly straightforward it has i number of decisions (W_i) and $i+2$ number of parameters (g_i, I_0 and F_i). It would be preferable if we can get the equation down to 1 decision.

To do this we examine the withdrawal amounts of each year. The goal of the retiree is to preserve his lifestyle. Ideally the amount of money withdrawn can purchase the same lifestyle each year. However the amount of money needed for a given lifestyle changes each year. The measurement of this change is captured in the inflationary index. Representing inflation as inf , we can calculate the relationship between successive withdrawals as:

$$W_{i+1} = W_i(1 + inf_{i+1}),$$

W_0 is defined as the initial withdrawal amount.

With some manipulation the equation becomes

$$F_i = ((I_0(1 + g_1) - W_0(1 + inf_1))(1 + g_2) - W_0(1 + inf_1)(1 + inf_2)) \dots W_{i-1}$$

Dividing both sides of the equation by the initial amount gives

$$F_i = (((1 + g_1) - \frac{W_0}{I_0}(1 + inf_1))(1 + g_2) - \frac{W_0}{I_0}(1 + inf_1)(1 + inf_2)) \dots \frac{W_0}{I_0} \prod_{k=1}^{i-1} (1 + inf_k)$$

We now have the equation in a state where the CWR for the initial amount can be explicitly solved. For a retirement over period r , the following equation holds.

$$\frac{W_0}{I_0} = \frac{\prod_{i=1}^{r-1} (1 + g_i) + F_r / I_0}{\sum_{i=1}^{r-2} [\prod_{j=1}^i (1 + inf_j) \prod_{j=i+1}^{r-2} (1 + g_j)] + \prod_{i=1}^{r-1} (1 + inf_i)}$$

It is important that in each retirement scenario the theoretical retiree does not use information from future years. This is important because this is the case in actual retirement. Therefore withdrawal is done at the end of each year because the inflationary adjustment will not be known until that point

We choose to interpret this process as the retiree has a certain cash reserve set aside from which they draw their day to day livings. At the end of the year the retiree then restores that amount with a withdrawal from their assets.

The advantage of the first is it takes into account the day to day living expenses of the client. The disadvantage is that the client has to have at least a year's budget in cash reserves, a fact made harder by the client not knowing what the inflationary rate is for that year. This can be mitigated by the observation that the inflation rate is unlikely to be 100 percent (in which case the client will notice and make appropriate modifications with their planner) so putting aside two years of expected withdrawals before retirement in cash will help.

Thus we have the concept of the *initial withdrawal rate (IWR)* equation. The IWR relates the approximate yearly buying power of a client for r years over a given set of growth and inflation.

$$IWR = \frac{\prod_{i=1}^{r-1} (1 + g_i) + F_r / I_0}{\sum_{i=1}^{r-2} [\prod_{j=1}^i (1 + inf_j) \prod_{j=i+1}^{r-2} (1 + g_j)] + \prod_{i=1}^{r-1} (1 + inf_i)}$$

Thus for the theoretical retiree in with saving I that wants to end with savings F can calculate their initial withdrawal rate if they know the next r years of growth and inflation.

A specific type of IWR is when F is set to 0. In that case the amount is called the *Safe Withdrawal Rate* as the retiree will just run out of money at the end of their retirement.

The problem with this approach is the inability to predict a successful retirement until all the growth and inflations are known. However there is strong evidence that the equation strongly depends on the interplay between the first few years of growth and inflation. If a reasonable prediction of whether or not a successful retirement can be reached with less than the full retirement period then considerable value could be derived.