Asset Allocation with Annuities for Retirement Income Management

April 2005

Paul D. Kaplan, Ph.D., CFA
Vice President, Quantitative Research
Morningstar, Inc.
225 West Wacker Drive
Chicago, IL 60606

1. Introduction

At, in, or near retirement, investors need to make specific decisions about how they are going to use their savings to generate income during retirement. These decisions include deciding how much to invest in annuities, how to invest non-annuitized assets, and how much to withdraw from the non-annuitized assets each year. Retired investors face the risk of running out of assets while they still need the income. Since the income needs usually remain until death, this risk is often called "longevity risk." Furthermore, many investors wish to leave some portion of their retirement accounts as part of their estates.

Retirement income management is choosing a combination of annuity, withdrawal, and investment strategies such that it is unlikely that the investor will run out of money before death while achieving the investor's financial goals. In making these decisions, the investor must manage both market risk and longevity risk.

In this paper, we explore retirement income solutions in a simple setting to illustrate the trade-offs that retired investors face regarding how much income they can generate, how much short-term risk they are exposed to, how large an estate they can expect to leave, and how likely they are to not run out of assets before dying (the "success" probability). We assume that at the beginning of the retirement period, the investor chooses how much annual income to generate (in real dollars), how much to invest in single premium immediate annuities, and what asset mix to use to manage non-annuitized assets.

We use a Monte Carlo simulation approach to solve the model for the probability of not running out of money before death with various combinations of real income, annuitization, and asset allocation. Our approach differs from many other simulation models in the way that we treat the level of income to generate. The usual approach is to fix the income level and solve for the probability of running out of money. In our approach, rather than picking an income level rate a priori, we find the trade-off curve between the income level and success probability. We then compare the trade-off curves across different combinations of annuitization and asset mixes to see how the choice of income level affects the highest achievable success probability and the combination of annuitization and asset mix needed to achieve that success probability.

The type of annuity that we explore pays a fixed nominal amount of money each year to the owner until death; so inflation causes the real value of the annuity payout to change over time. Hence, we need to include a model of inflation in our analysis. We construct three different inflation models, all with the same expected inflation rate, and find the relationship between the real income goal and the highest achievable success probability under each model. Our analysis shows that how we choose to model the inflation process can have a substantial impact on this relationship and on optimal investment and annuitization decisions.

Success probability may not be the only criterion for selecting a withdrawal and investment strategy. Potential future wealth, which becomes the estate value at death, may also be a criterion. So for a given income level, we also calculate median estate values as well as the success

probability for various combinations of annuitization and asset mixes to find the trade-off between success probability and potential estate size under one of our inflation models.

The remainder of this paper is organized as follows: Section 2 reviews the basic withdrawal model that is commonly used to address the issues of retirement income management with simulation in which the income level is a given. This provides the background for the literature review in Section 3. Section 4 presents our basic model without annuities in which we find the trade-off curves between the income level and the probability of not running out of money before death. Section 5 presents the three inflation models. Section 6 introduces annuities into our model and shows the impact of annuitization on the trade-off between income and success probability under each of the inflation models. It also shows the trade-off between success probability and potential estate size for a given income level under one of the inflation models. Section 7 discusses extensions to the model. Section 8 summarizes our approach and findings.

2. The Simulation Approach to Withdrawal Models

The typical approach to analyzing the problem of generating retirement income from a pool of invested assets is to run a series of simulations either by using historical return data or simulating returns with a parametric model.[1] Figure 1 shows the structure of a typical withdrawal experiment. We start with a portfolio of a given value, say $500,000. The money in the portfolio is invested in some assets. The assets generate returns over a period of time, say a year. For example, if the assets generate a return of 5%, the portfolio is worth $525,000 before withdrawals. Then a withdrawal is made. (In reality, withdrawals have to be made more frequently, but for the purpose of modeling, we assume that withdrawals are made at the end of each year.) If the withdrawal was $50,000 for the year, the portfolio drops in value to $475,000.

If the person now dies and the portfolio value is at least as large as the estate goal, the experiment ends as a success. If the person dies and the portfolio value is less than the estate goal, the experiment ends as a failure. If the person is still alive, the portfolio is rebalanced to the target asset allocation and the experiment continues.

Each experiment is performed with a different sequence of asset returns. The result of each experiment (success or failure) is recorded. The percentage of experiments that end in failure measures the risk of the combination of withdrawal policy, investment strategy, and estate goal. The more independent experiments that are performed, the higher is the accuracy of the measured risk.

Generating income while avoiding running out of money may not be the only criterion for selecting an investment strategy. The investor might also be interested in what sort of portfolio balance he or she can expect to have over time. A larger portfolio balance provides the opportunity to either

---

[1]   For example, Cooley, Hubbard, and Waltz [1998] and Jarrett and Stringfellow [2000] use historical return data, while Milevsky, Ho, and Robinson [1997], McCabe [1998], and Pye [2000] simulate asset class returns with a parametric models.

generate greater income, leave a larger estate, or a combination of both. The investor may also be concerned with the possibility of facing unplanned major expenses at some point in the future that would require dipping into the retirement income account. For all of these reasons, the investor might be concerned not only with the account balance at death, but at any point in time.

Given the withdrawal rate and investment returns, we can depict the distribution of the account balance over time graphically. Figure 2 shows various percentiles of a simulated account balance for someone who starts with $500,000 at age 65 and withdraws $25,000 per year and follows a particular investment strategy. The figure shows that in this example, there is a 10 percent chance that the investor will run out of money by age 89 and a 25 percent chance that the investor will run out of money by age 97. Should the investor die at an earlier age or if the investments have better performance, he or she will leave the account balance as part of an estate.

3.   Literature Review

Over the past decade, a number of papers have been published on the role of asset allocation and annuitization in retirement income management. These studies differ in approach and emphasis.

Some studies focus on sustainable withdrawal rates for portfolios of stocks, bonds, and cash using historical asset class return data. For example, Cooley, Hubbard, and Waltz [1998] find the number of overlapping fixed length periods for which a given inflation-adjusted withdrawal rate would have been sustainable for a variety of period lengths, withdrawal rates, and asset mixes over the period 1926-1995. Using similar historical data on asset class returns for the period 1926-1998 and a similar set of asset mixes, Jarrett and Stringfellow [2000] calculate the highest withdrawal rate that would have been sustainable in each and every overlapping fixed length period for each fixed period length that they examined.

Several studies examine the relationship between the withdrawal rates and success probabilities for different asset mixes using a parametric model of real asset returns. McCabe [1998] uses the method described in Section 2 above for fixed time horizons with a lognormal model of asset class returns. Pye [2000] performs a similar exercise, but varies the withdrawal during the simulation based on account balance and remaining length of time remaining so as to increase the probability of not running out of money before the end of the period.

Using a multivariate lognormal model of real asset class returns, Milevsky, Ho, and Roberson [1997] use Monte Carlo simulation to find the probability of running out of money before death for a given real withdrawal rate when the time of death is uncertain with asset mixes that vary from 0 to 100 percent equity. To model the uncertainty of time of death, they first calculate the probability of running out of money for each year using the technique described in Section 2 above. They also calculate the probability of the investor dying in each year based on the person's age and gender using a standard mortality table. They calculate the overall probability of running out of money as the weighted average of the probability of running out of money in each year, using the probability of death in each year as the weights. As discussed in Section 4 below, we use essentially the same

approaches as Milevsky, Ho, and Roberson to modeling asset class returns and handling the uncertainty of the year of death in our model.

Monte Carlo simulation is inherently a computationally intensive method of approximation. So some researchers have developed analytical approximations for the probability of running (or not running) out of money under simplified assumptions. McCabe [1999] presents an approximation for the probability of never running out of money over an infinite horizon when withdrawing a constant amount, given the mean and standard deviation of the return on the invested portfolio. Milevsky and Robinson [2000] present a different solution to this problem that is exact if the returns on the invested portfolio follow a lognormal distribution. Milevsky and Robinson also present an analytical approximation for the probability of running out of money before death when the time of death is uncertain. However, the numerical analysis that is required to carry out calculations may be computationally intensive.

A number of papers present formal solutions to optimization problems in which the investor selects an asset mix (and in some cases a level of annuitization) so as to maximize either an expected utility function or the probability of not running out of money before death. For example, Young [2004] derives the optimal allocation between a risk-free and a risky asset for an investor who needs a constant real level of income, has a constant probability of death, and wants to minimize the probability of outliving wealth. In Young's model, the optimal proportion for the risky asset is a decreasing function of wealth, and hence changes over time.

Milevsky, Moore, and Young [2004] introduce a constant dollar single premium immediate annuity into Young's model. In their model, due to the binary nature of the investor's goal, the optimal solution is to do no annuitization until the investor can generate the desired level of income entirely with annuities. At that point, the investor sells off all assets and purchases the annuity.

In an expected utility model, typically the level of consumption is treated as variable rather than as a fixed income level. Hence, papers such as Milevsky and Young [2003] examine the optimal strategy for consumption as well as for annuitization and asset allocation.

One of Milevsky and Young's findings is that if investors can annuitize a fraction of wealth at distinct points in time, they will initially annuitize a lump-sum and then keep buying annuities in order to keep the ratio of annuities held to non-annuitized wealth at least as great as an optimal fixed ratio. This result is in sharp contrast to the optimal level of annuitization being all or nothing as it is in Milevsky, Moore, and Young's model.

Chen and Milevsky [2003] present an expected utility maximizing model in which investors allocate their wealth to conventional equities, conventional fixed income, fixed annuities, and an equity-based variable annuity. However, rather than using a multi-period model of investing and consumption, they use a single period model. In the single future period, the investor is either alive or dead. In the living state, all four types of assets have a payoff. In the dead state, only the non-annuity assets have a payoff. So the optimal allocation between annuity and non-annuity assets depends on the strength of the estate motive. Since it is a single period model, it does not address

the generation of income over time. However, the authors indicate that they will present a multi-period in a future paper that presumably will address income generation.

In all of the papers discussed above, the fixed annuity pays a constant real amount. To our knowledge, the existing literature does not address the fact that most fixed annuities have constant nominal payouts rather than real payouts. The model presented in this paper may be the first to address the issue of using nominal fixed annuities as part of a strategy to generate a constant real level of income.

4. A Model Without Annuities

We assume that all non-annuitized assets are invested in diversified portfolios of stocks, bonds, and cash. We model the real returns on these three broad asset classes as following a multivariate lognormal distribution. Our assumptions on expected returns, standard deviations, and correlations are based on Morningstar Associates' model of real asset class returns. Figure 3 shows the expected real compound returns and standard deviations implied by this model.

Morningstar Associates has a set of asset allocation policy mixes that vary from 0 to 98 percent stocks as shown in Figure 4. We select 21 of these asset mixes for our drawdown analysis; namely those with stock allocations of 0, 5, 10, …90, 94, and 98 percent. Figure 3 shows the implied expected real compound annual returns and standard deviations on these 21 asset mixes along with these parameters for the full set of policy asset mixes. As this figure shows, the selected mixes cover a full range of points along of risk-return spectrum.

We assume that the investor withdraws a constant real dollar amount from the invested assets each year. If there is inflation, the nominal value of the withdrawn amount rises in lockstep with the price index. Note that since we model both asset class returns and withdrawal rates in real terms, we do not need an explicit model of inflation when there are no annuities. (In next section, when we introduce annuities that pay a fixed nominal amount, we introduce explicit models of inflation.)

We perform 1,000 Monte Carlo experiments. For each experiment at each horizon, we find the highest withdrawal rate that would leave the investor with a balance of zero.[2] At each horizon, we sort these withdrawal rates across the experiments to find the relationship between withdrawal rate and success probability at the horizon in question. For example, if at the 20-year horizon, the 10th percentile of withdrawal rates were 5.8 percent of the initial balance, we would conclude that the probability of not running out of money in 20 years is least 90 percent if the withdrawal rate were no more than 5.8 percent of initial wealth. The uneven curve in Figure 5 shows the relationship between the withdrawal rate and the success probability at the 20-year horizon for the 40 percent stock asset mix that we obtain by this method.[3]

---

[2]   In a similar fashion, Jarrett and Stringfellow [2000] calculate sustainable withdrawal rates for overlapping historical periods.

[3]   In a similar fashion, McCabe [1998, 1999] presents results using graphs of functions that relate the withdrawal rate to the probability of success.

In theory, the relationship between the withdrawal rate and the success probability at each horizon should be a smooth curve. The unevenness of the relationship derived from the Monte Carlo simulation is due to approximation error inherent in the Monte Carlo method. Since the analysis that we do below requires smooth functional relationships between the withdrawal rate and the success probability, at each horizon, we use a smooth curve through selected points to represent the entire relationship. As illustrated in Figure 5, the smooth curves that we form are very close to the uneven curves that we derive them from.

Since the year of death is uncertain, we take a weighted average of the success probabilities at the various horizons (as determined by the smoothed functions) as the overall success probability for a given withdrawal rate. The weight for each horizon is our estimate of the probability that the investor will die in that particular year. We derive our death-year probability estimates from the mortality rates published by the Society of Actuaries.[4] Figure 6 presents these gender specific mortality rates.

For our example of retirement income management, we use a 65-year old male who has just retired. Using the mortality rates for males shown in Figure 6, we derive the death-year probabilities for a 65-year old male shown in Figure 7. From these probabilities, we find that while the life expectancy of this person is 84 years, there is nearly a 50 percent probability that he will live beyond his life expectancy. This is why we think that it is important to use a full probabilistic model of the horizon rather than focusing just on the life expectancy when accessing retirement income strategies.

The weighted-average success probability associated with each withdrawal rate defines a functional relationship between the withdrawal rate and the overall success probability. We find the withdrawal rates associated with a selected set of probabilities and fit a smooth curve through resulting withdrawal rate/success probability pairs. To see how the choice of asset mix affects the resulting relationship, we repeat these calculations for each of our selected asset mixes.

Figure 8 shows the resulting relationships between withdrawal rates and success probabilities for each of the 21 asset mixes that we selected. The right edge of the resulting figure defines an efficient frontier between withdrawal rates and success probabilities, assuming that the investor would select the asset mix to maximize the success probability. Figure 9 shows the resulting equity allocations. Note that the equity allocation increases as the withdrawal rate increases and the success probability falls.

Figure 8 shows that in this example, to keep the success probability in the 90-95 percent range; the withdrawal rate should not exceed 5 percent of the initial portfolio value. At this withdrawal rate, the highest success probability rate is 92.6 percent. Figure 9 shows that this occurs at the asset mix with an equity allocation of 40 percent.

---

[4]   See Johansen [1998].

5.   Modeling Inflation

Since the annuity pays a fixed nominal amount each year, we need a model of the annual inflation rate to express these payments in real terms. To capture the uncertainty of inflation, the inflation model should generate time series of annual inflation rates that differ between experiments.

Annual data on U.S. inflation rates exhibits serial correlation. To explore the impact that the serial correlation of inflation might have on the model, we use three alternative models of the annual inflation rate. Let

π_t = the continuous annual inflation rate for year t
π_M = the continuous long-term annual inflation rate

For our first model, we assume that the inflation rate is constant:

$$\pi_t = \pi_M$$

For our second model, we assume that there is first order serial correlation in π_t:

$$\pi_t = 0.1145\pi_M + 0.8855\pi_{t-1} + 0.0128z_t$$

where z_t is serially independent white noise that follows a standard normal distribution. We estimated the coefficients on π_{t-1} and z_t using historical inflation data for the period 1978-2003.[5] The coefficient on π_M is one minus the coefficient on π_{t-1} so that π_t will tend towards π_M. We also estimated the correlation of z_t with each of the real returns on the asset classes using historical data over the same period.

For our third model, we assume that there is second order serial correlation in π_t:

$$\pi_t = 0.2222\pi_M + 1.3876\pi_{t-1} - 0.6098\pi_{t-2} + 0.0086z_t$$

We estimated the coefficients on π_{t-1}, π_{t-2}, and z_t using the same historical data as we did for the one-lag model. The coefficient on π_M is one minus the sum of coefficients on π_{t-1} and π_{t-2} so that π_t will tend towards π_M. We also estimated the correlation of z_t with each of the real returns on the asset classes using historical data over the same period.

In our examples, we assume an initial inflation rate of 2.5 percent per year. We also assume that the autoregressive process causes the annual inflation rate to mean revert to 2.5 percent. Hence we set π_{-1}, π_0, and π_M all to ln(1.025) in all of the models.

---

[5]   The monthly inflation rates that Morningstar Associates uses to represent inflation over the period 1/78-12/03, linked into calendar years. The underlying measure of inflation is the rate of change of the seasonally adjusted Consumer Price Index published by the Bureau of Labor Statistics.

To each of the 1,000 experiments that we ran in Section 4, we add simulated values of z_t that are correlated with the real returns on the asset classes for the one-lag and two-lag models, the correlation coefficients depending on the choice of the inflation model. Under each inflation model, we generate a time series of values of π_t for each experiment. We compound the inflation rates to generate a time series for the price index in each experiment. We do this by setting the price index in year 0, P_0, to 1 and then calculating the price index in each year t, P_t as follows:

$$P_t = P_{t-1} \exp(\pi_t)$$

The higher the level of the price index, the lower will be the value of the annuity to the investor. To see the potential impact that each inflation model might have on the income generation model, we calculate the average level of the price index across the 1,000 simulation experiments for each future year under each model. Figure 10 presents the results. As the figure shows, constant inflation generates the lowest values for the price index, the two-lag model generates slightly higher values, and the one-lag model generates significantly higher values.

The models have different price index averages because they differ in inflation volatility. The mathematics of compounding is such that the more volatile a rate-of-change series is, the greater the compound value, all other things being equal. By definition, the constant inflation model has no volatility; hence it produces the lowest average.

The other two inflation models produce volatile inflation rates. However, they do so in different manners. To see the difference between the two models, we calculate their autocorrelation functions. Figure 11 presents the results. As the figure shows, the one-lag model produces long monotonic trends while the two-lag model produces shorter cyclical trends. The longer trends of the one-lag model result in inflation rates that have greater variation around the mean than those in the two-lag model.

To see the differences in the volatilities between the two models, we calculate their respective forecast volatilities. Figure 12 presents the results. This figure shows that forecast the volatility of the two-lag model is lower than that of the one-lag model at all horizons and significantly so within ten years. Hence, the simulated inflation rates generated by the one-lag model are more volatile than those generated by the two-lag model.

6.   Adding Annuities to the Model

To explore the potential role of fixed annuities in retirement income management, we modify our model by allowing the investor to allocate a portion of initial wealth to a single premium immediate annuity that pays a fixed nominal amount each year until death and has no beneficiaries. For simplicity, we only allow the annuitization decision to be made at the beginning of the retirement period. As with almost all actual annuity contracts, the decision to annuitize is irreversible.

Fixed annuity prices are based on mortality rates and hence on age and gender. In our example of a 65-year old male, we assume that $7.50 is received each year of life for every $100 paid at retirement.

Under each of our three inflation models, we calculate values of the price index from the simulated inflation rates in each experiment and divide the resulting simulated price index values into the nominal annual annuity payout to find the real annuity payout. Thus, while the nominal annuity payout remains constant, the real annuity payout varies across time and under our second and third inflation models, across experiments as well. Since the long-term inflation rate is positive, the real annuity payout tends to decline over time.

To see the effect on the trade-off between real income level and success probability when fixed nominal annuities are available, we again do a Monte Carlo simulation for a 65-year old male. We assume that this person starts with a balance of $500,000 and seeks to generate a constant level of income per year in real terms by dividing the initial balance between single premium immediate annuities that yield 7.5 percent per year on the initial purchase and a portfolio of risky investments assets. In years in which the annuity payments fall short of the desired income level, the investor withdraws money from the invested assets to make up the shortfall. In years in which the annuity payments exceed the desired income level, the investor adds the excess to the pool of invested assets. At the end of each year, after either withdrawing or adding funds, the investor rebalances the invested assets back to a fixed asset mix.

We use 20 levels of annuitization (0, 5, 10, …, 95 percent of initial wealth) and our 21 selected asset mixes to form 420 combinations of annuitization and equity allocation. For each combination, we find the relationship between the income level (expressed as a percentage of initial wealth) and the success probability using the same technique that we use for the model without annuities. For each income level, we find the combination of annuitization and asset mix that provides the highest success probability. We repeat the entire exercise using each of the three inflation models.

Figure 13 illustrates that under the one-lag inflation model, for the 5 percent income level, using 70 percent of the initial balance to purchase annuities and investing the remaining 30 percent of the initial balance in a portfolio that is 70 percent equity provides the highest success probability. Figure 14 presents the resulting trade-offs between income levels and maximum success probabilities without annuities, and with annuities under each of the three inflation models.

As Figure 14 illustrates, the improvement in success probability depends largely on the choice of the inflation model. Under all three models in this example, the largest improvement occurs when the target income level is about 5 percent of initial wealth. The following table summaries the results for this real income target:

Model	No Annuities	Constant Inflation	One-Lag Inflation	Two-Lag Inflation
Annuitization	0%	95%	70%	85%
Equity Allocation of Non-Annuitized Assets	40%	30%	70%	70%
Success Probability	92.56%	99.52%	94.66%	97.32%

These results are consistent with our earlier remark that inflation models that result in lower values of the price index will make annuities more beneficial for the investor. The most striking result is that under the constant inflation model, the real level of income can be almost fully guaranteed using a nominal annuity. This is because in the early years, when the real level of income generated by the annuity exceeds the target, the excess income is added to the invested portfolio. In later years, when the real income from the annuity falls short of the target, the investor makes up the gap by making withdrawals from the invested portfolio.[6]

In fact, if the real income goal is not too high and inflation is perfectly forecastable, it is possible to completely eliminate longevity risk using annuities. Initially, the investor purchases enough annuities to both meet the first year's income goal and to provide the funds for future annuity purchases that will be needed to maintain the real level of income. Subsequently, in each year, the investor spends the inflation-adjusted income goal and uses the excess to buy additional annuities. All other things being equal, the annuity price should fall as the investor ages, making this strategy more economical over time. Any funds not needed for annuities can be invested in conventional assets to build an estate. In general, it should be possible to create a dynamic strategy of annuitization and conventional assets for retirement income management and estate planning using nominal annuities when inflation is uncertain; but this is a topic for future research.

Annuitization and asset allocation decisions impact financial outcomes in addition to the probability of not running out of money before death. In many cases, the annuitization decision is largely one of securing a level of retirement income versus maintaining control over one's assets, especially for the purpose of leaving an estate. To illustrate this trade-off, we repeat the Monte Carlo simulation with the one-lag inflation model for the 65-year male who has access to a fixed nominal annuity that yields 7.5 percent. For each of the 420 combinations of annuitization and equity allocation, we calculate the probability of failure (i.e., running out of funds before death) and median estate value.[7] Figure 15 presents the resulting combinations of probability of failure and median estate value.

The left edge of the cluster of points in Figure 15 traces out an "efficient frontier" of probability of failure versus median estate value. In Figure 16, we plot just the points on the most efficient points,

---

[6]   To actually implement this strategy, the investor would have to use a taxable investment account. The model presented here ignores taxes.

[7]   The probability of failure is simply 100 percent minus the success probability. We plot the probability of failure rather than the success probability in Figures 13 and 14 to present the trade-offs between risk and reward in the usual manner.

with and without annuities.[8] The gap between the two sets of points shows how much annuities can improve the trade-off between risk and reward in this setting.

For a given combination of annuity allocation, θ, and equity allocation of non-annuitized assets, x, the effective allocation to stocks is $(1-\theta)x$ and the effective allocation to conventional fixed income is $(1-\theta)(1-x)$. In Figure 17, we plot the effective allocations to annuities, conventional fixed income, and stocks implied by the combinations of θ and x depicted in Figure 16 that make up the efficient frontier with annuities. Figure 17 is a simplex diagram for allocations to three asset classes. In this diagram, a 100 percent allocation to an asset class is represented by a point at the corner of the triangle labeled with the name of the asset class. A point on the side of the triangle represents an allocation that has no allocation to the asset class named on the opposite corner, but is entirely made up of the asset classes named at the corners that are the endpoints of the side in question. In general, any point on the triangle represents an allocation to the three asset classes; the closer the point to a particular corner, the higher the allocation to the asset class the corner represents.

The pattern of points in Figure 17 suggests that the efficient frontier consists of three segments:

- High stock allocations with no allocation to annuities
- A segment from the zero annuity segment to the point of minimum probability of failure
- A segment from the point of minimum probability of failure to 100 percent annuities[9]

Note that along the last segment, as with any straight line on the triangle in Figure 17 that goes through the 100 percent annuities point, x is a constant.

The adjacent line segments shown in Figure 17 are drawn as described above. To construct these line segments, we need to identify two combinations of θ and x:

- The least risky efficient asset mix that has no allocation to annuities
- The combination that minimizes the probability of failure

To see if the adjacent line segments in Figure 17 provide a good approximation of the allocations of the efficient frontier, we map 99 points along the segments, equally spaced by stock allocation, into combinations of θ and x (x = stock allocation / $(1-\theta)$) and calculate the probability of failure and median estate value for each combination. The results are shown as the continuous curve in Figure 16. As the figure shows, the approximated efficient frontier falls almost exactly on the efficient points.

In Figure 18, we show the effective asset mixes between stocks, bonds, cash, and annuities for the 99 combinations of annuitization and equity allocation depicted by the curve in Figure 16 and by the

---

[8]   Strictly speaking, the points below where probability of failure is minimized are not efficient. However, we include them in our "efficient frontier" here because the lower negatively sloped segment is often shown in efficient frontier diagrams drawn in the standard deviation-expected value space of traditional portfolio theory.

[9]   See previous note.

adjacent line segments in Figure 17. We generated this chart by applying the stock, bond, cash mixes shown in Figure 4 to the non-annuitized portion of the 99 portfolios.

7. Extensions to the Model

Our model can be readily extended in a number of ways to more realistically reflect the actual circumstances of retirees. First of all, retirees typically have sources of income besides their retirement nest eggs. These additional income sources may include Social Security benefits, payouts from defined benefit plans, income from part-time jobs held during retirement, and income generated by real estate investments. These other sources of income can be combined with annuity payouts and withdrawals from invested assets to create a constant level of real income. Other cash flows, positive or negative, that may take place during retirement can also be incorporated into the analysis.

The model can be extended to the case of two persons retiring together, using a single pool of money to generate income, as is typically the case with retired married couples. In this version of the model, we model each person's probability of death each year based on age and gender and assume that the death of each person is an independent event. Hence in each year, there are four scenarios, each with an associated probability; namely,

- Both people remain alive
- The first person dies, survived by the second person
- The second person dies, survived by the first person
- Both people die

We assume that when one person survives the other, a fixed fraction of the joint income goal becomes the income goal of the survivor. Once both people die, any remaining balance in the invested portfolio becomes the estate.

The single premium immediate annuity used in the single person model is replaced by a single premium immediate joint-and-survivor annuity in the two-person model. In a joint-and-survivor annuity, the payout depends on which of the members of a couple are alive, the first person only, the second person only, or both. Payouts stop once both people die.

In the analysis presented here, we focus on the real income level, the probability of not running out of money, and estate size. Another factor that we can consider is the short-term risk of the investment portfolio. All other things being equal, an investor may want to avoid taking on short-term risk in the invested portfolio. We can incorporate such short-term risk aversion directly into our model by including a short-term risk measure, such as portfolio volatility, directly into the objective function when selecting the asset mix and the level of annuitization.

8. Summary

During retirement, both market and longevity risk must be managed by making choices regarding income generation, annuitization, investments, and estate planning. Simulation models can be useful to investors planning retirement income by portraying the trade-offs between income level, portfolio volatility, success probability, and estate size that are available through different combinations of asset mixes and annuities. In this paper, we considered what trade-offs are available through a set of asset mixes and single premium immediate annuities.

When there are no annuities in the model, there is no need to model inflation since all variables are in real terms. However, since fixed annuities pay a constant nominal amount, we need to model inflation to model their real payouts. We found that the role that annuities play in an optimized strategy depends largely on how we choose to model inflation.

For a given real income goal, there is an efficient trade-off between failure probability and median estate value. Points along this efficient frontier have different combinations of equities, conventional fixed income, and fixed annuities that we can depict graphically.

The model can be extended to consider additional income sources and other fixed cash flows, two-person mortality, and short-term risk aversion. With these extensions, the model can be used as a guide to help retired investors make informed decisions on income generation, annuitization, asset allocation, and estate planning.

Our analysis focused on strategies in which the investor makes a one-time annuitization decision and selects a single asset mix for non-annuitized assets. Further research needs to be done to explore dynamic annuitization and asset allocation strategies during retirement in the presence of uncertain inflation.

References

Chen, Peng and Moshe A. Milevsky, "Merging Asset Allocation and Longevity Insurance: An Optimal Perspective on Payout Annuities," *Journal of Financial Planning*, June 2003.

Cooley, Phillip, Carl Hubbard, and Daniel Waltz, "Retirement Savings: Choosing a Withdrawal Rate that is Sustainable," *AAII Journal*, February 1998.

Jarrett, Jaye and Tom Stringfellow, "Optimum Withdrawals from an Asset Pool," *Journal of Financial Planning*, January 2000.

Johansen, Robert J., "Annuity 2000 Mortality Tables," in Transactions*, Society of Actuaries*, 1998.

McCabe, Bernard J., "Monte Carlo Analysis of the Impact of Portfolio Volatility on After-Tax Wealth," *Journal of Private Portfolio Management*, Winter 1998.

McCabe, Bernard J., "Analytic Approximation for the Probability that a Portfolio Survives Forever," *Journal of Private Portfolio Management*, Spring 1999.

Milevsky, Moshe A., Kwok Ho, and Chris Robinson, "Asset Allocation via the Conditional First Exit Time or How to Avoid Outliving Your Money," *Review of Quantitative Finance and Accounting*, July 1997.

Milevsky, Moshe A., Kristen S. Moore, and Virginia R. Young, "Optimal Asset Allocation and Ruin-Minimization Annuitization Strategies," Individual Finance and Insurance Decisions Centre working paper, May 2004. Available at http://www.ifid.ca/pdf_workingpapers/WP2004May10.pdf.

Milevsky, Moshe A. and Virginia R. Young, "Annuitization and Asset Allocation," Individual Finance and Insurance Decisions Centre working paper, October 2003. Available at http://www.ifid.ca/pdf_workingpapers/WP2003OCT15.pdf.

Milevsky, Moshe A. and Chris Robinson, "Self-Annuitization and Ruin in Retirement," *North American Actuarial Journal*, October 2000.

Pye, Gordon B., "Sustainable Investment Withdrawals," Journal of Portfolio Management, Summer 2000.

Young, Virginia R., "Optimal Investment Strategy To Minimize The Probability Of Lifetime Ruin," *North American Actuarial Journal*, October 2004.
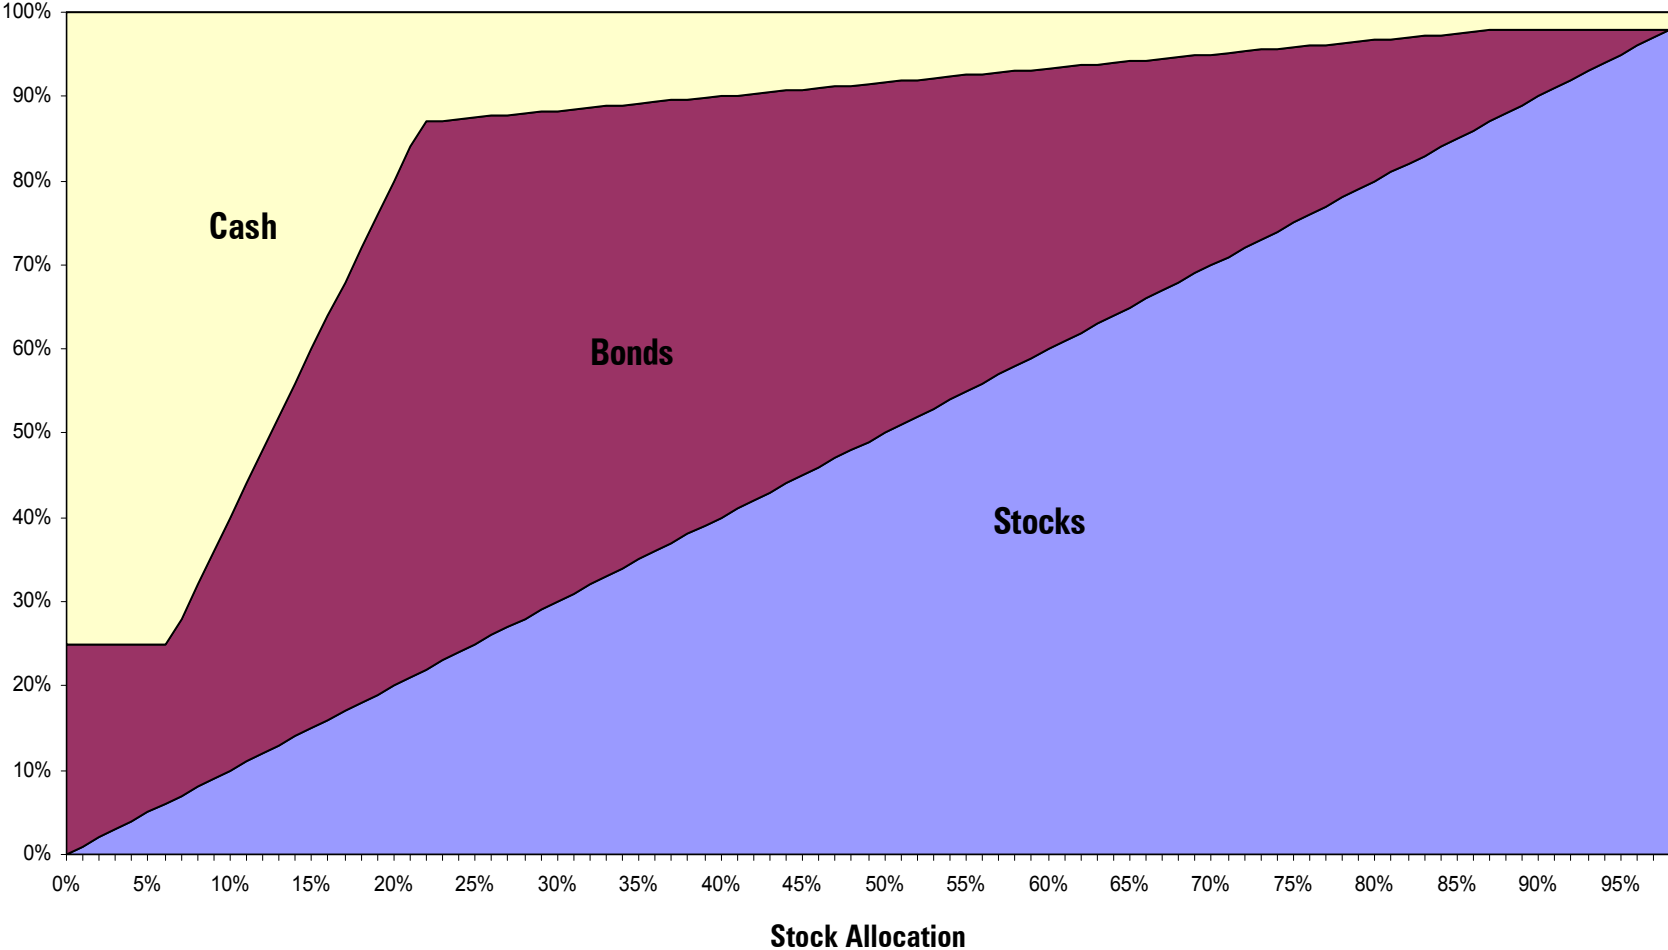
Figure 1: Simulation Using Withdrawal Experiments
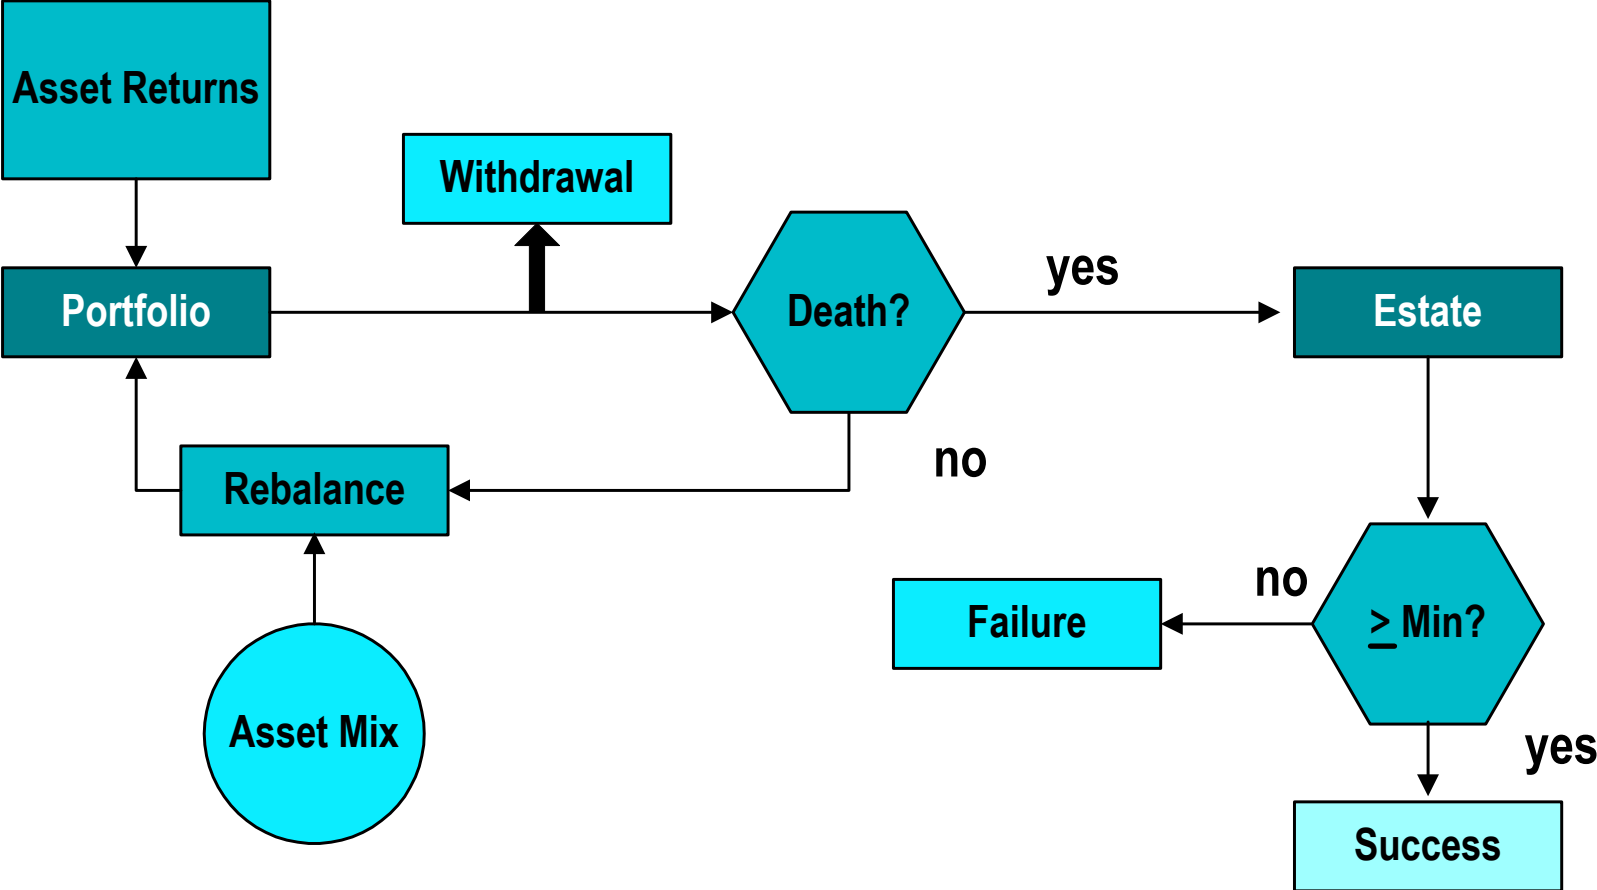
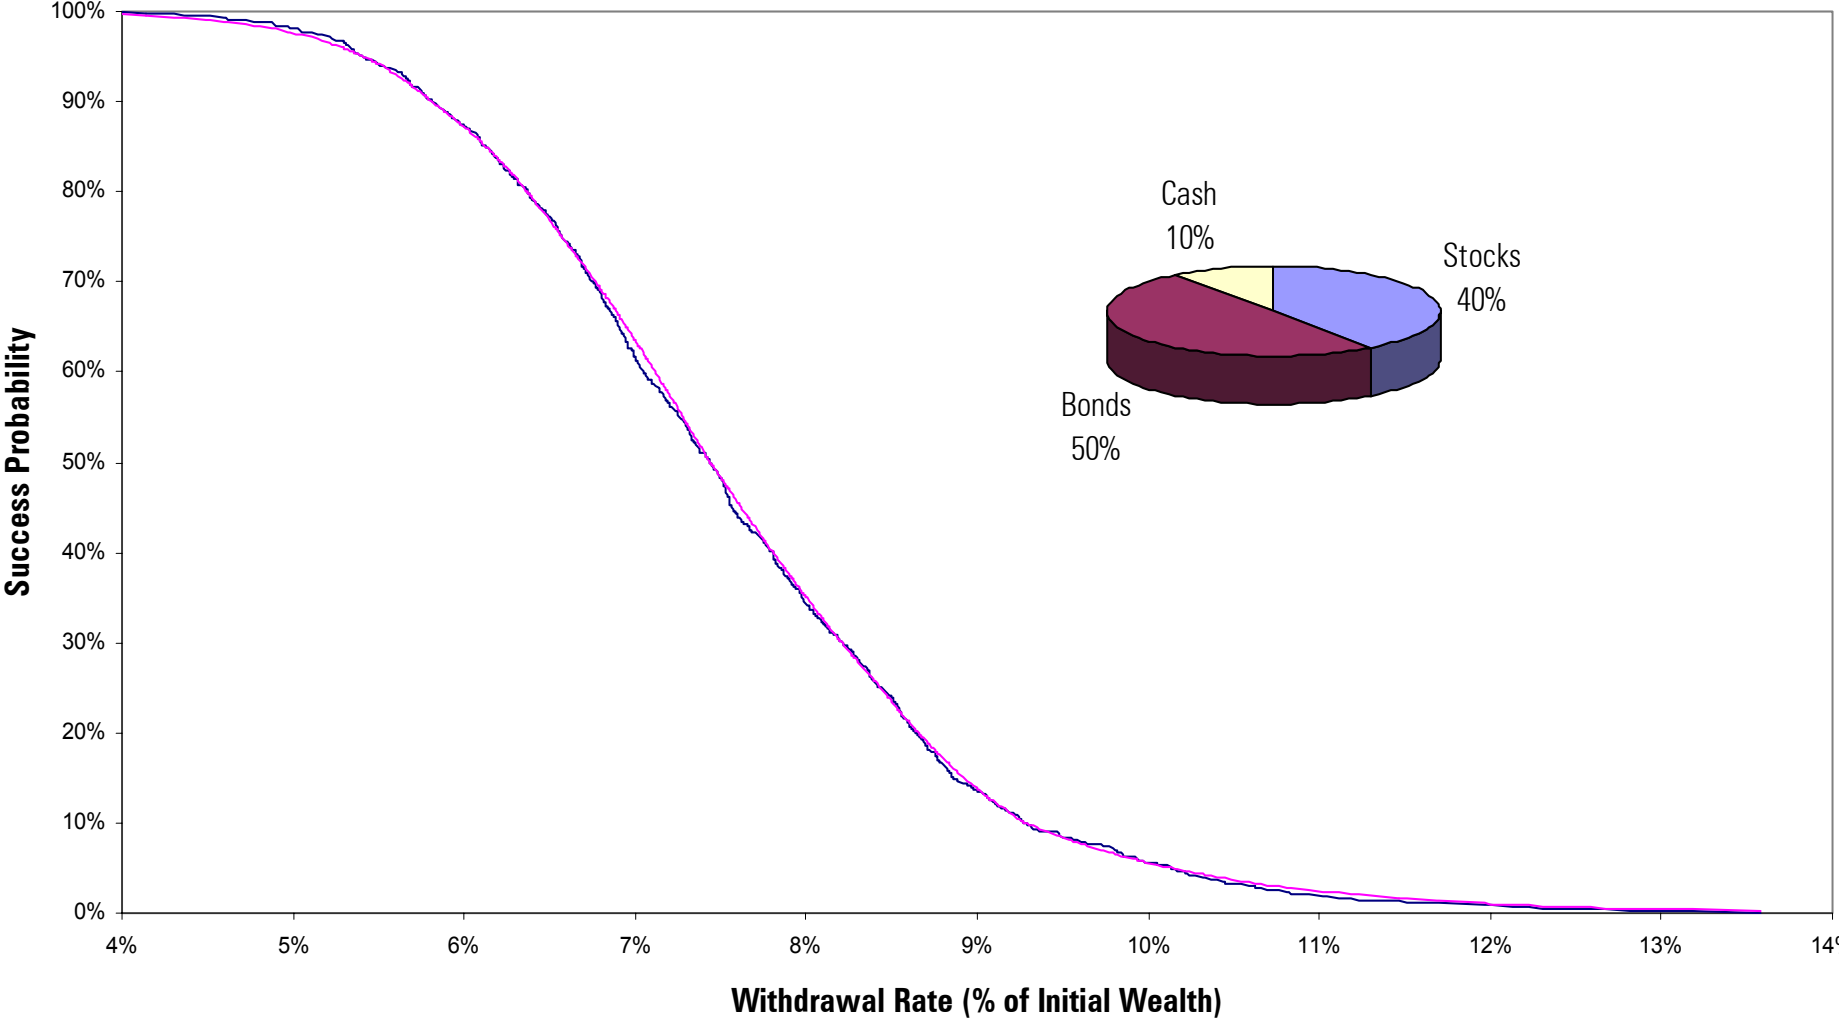**Figure 2: Percentiles of Simulated Balances at Various Ages**
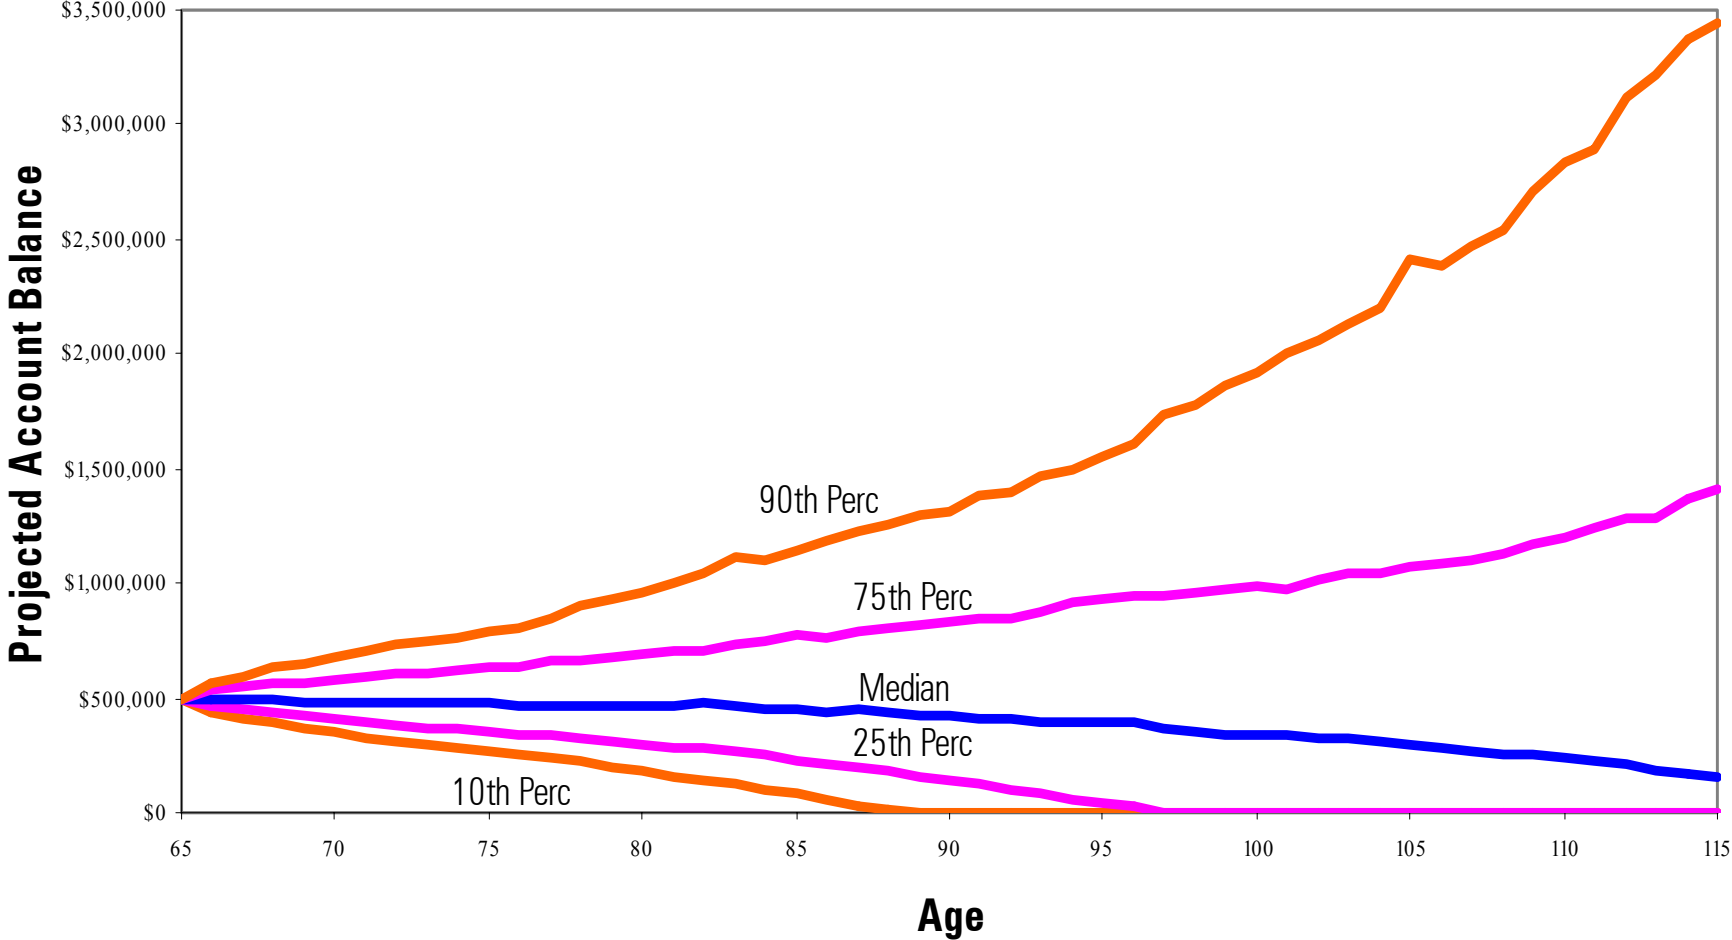
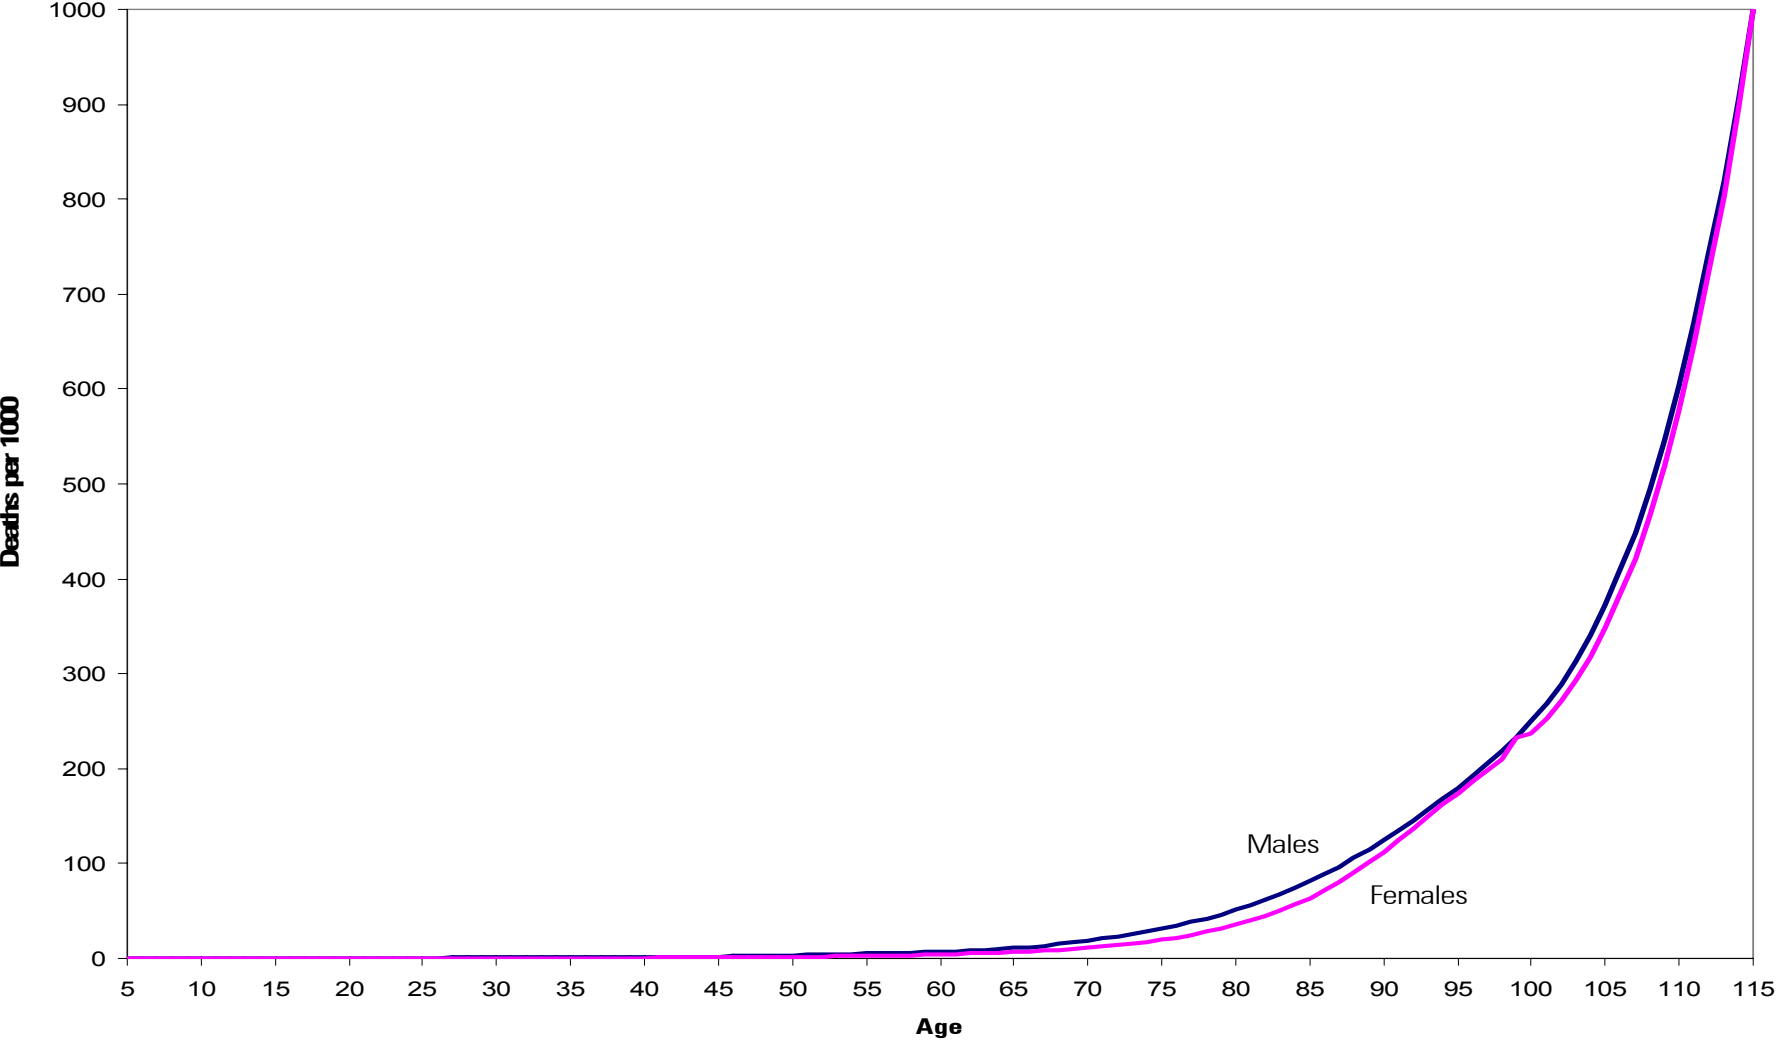# Figure 3: Real Return Assumptions on Asset Classes and Mixes
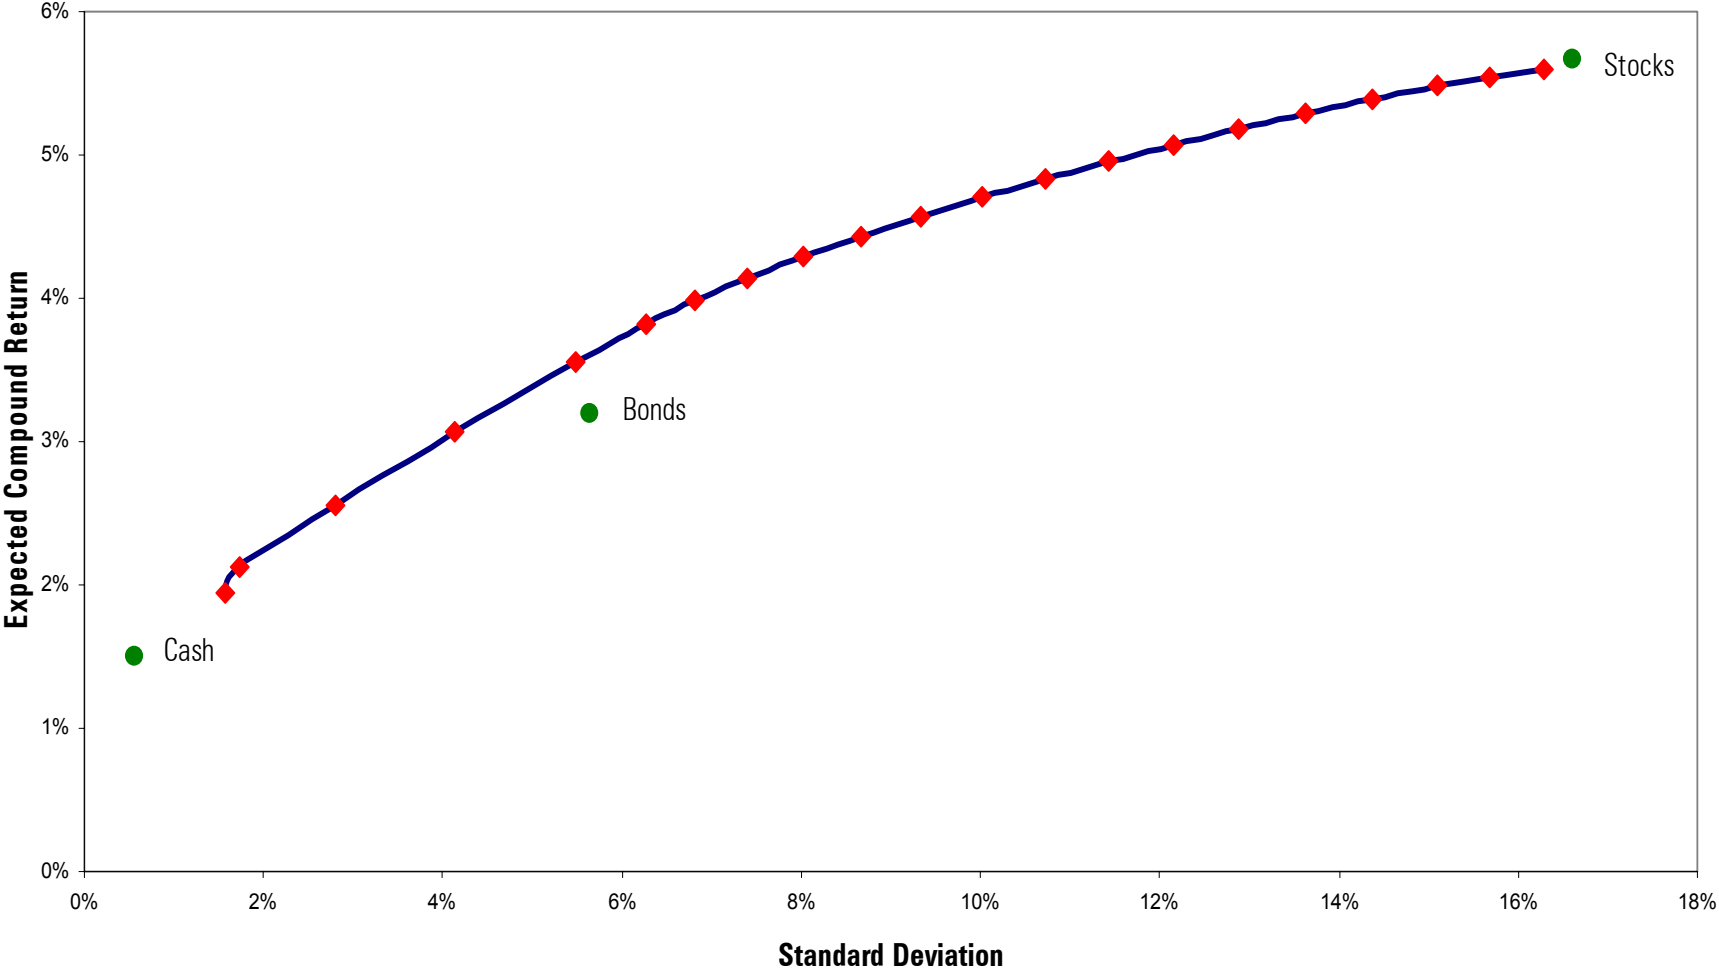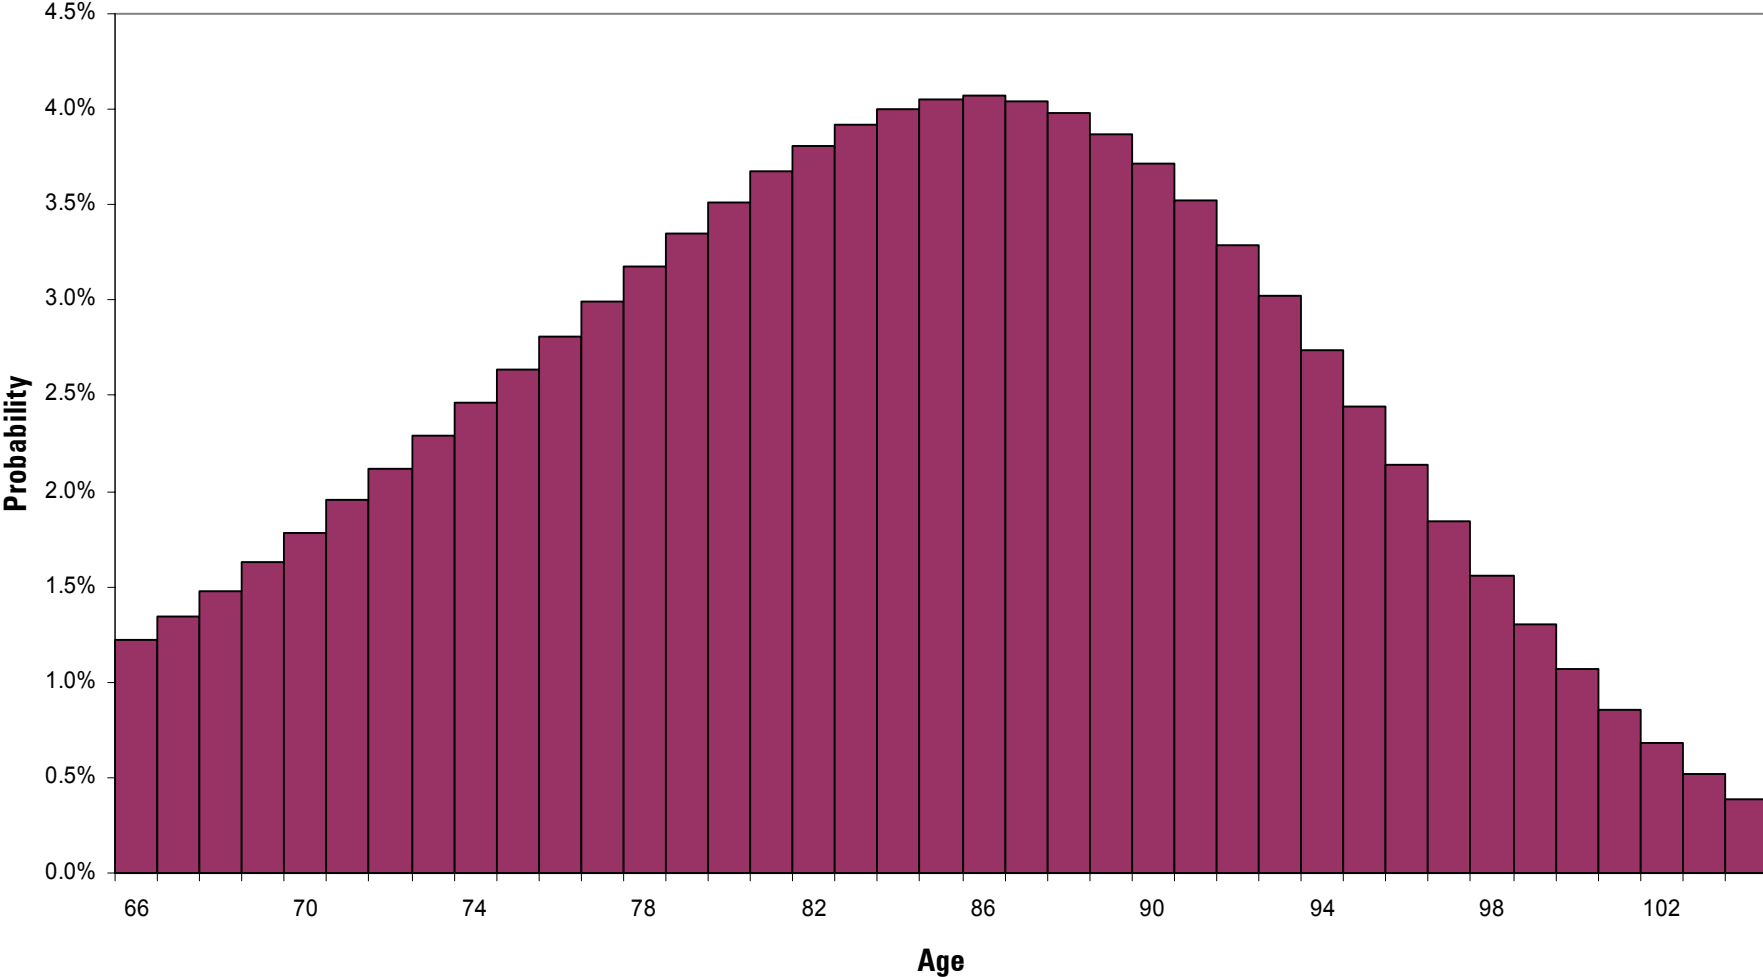
Figure 4: Asset Mixes

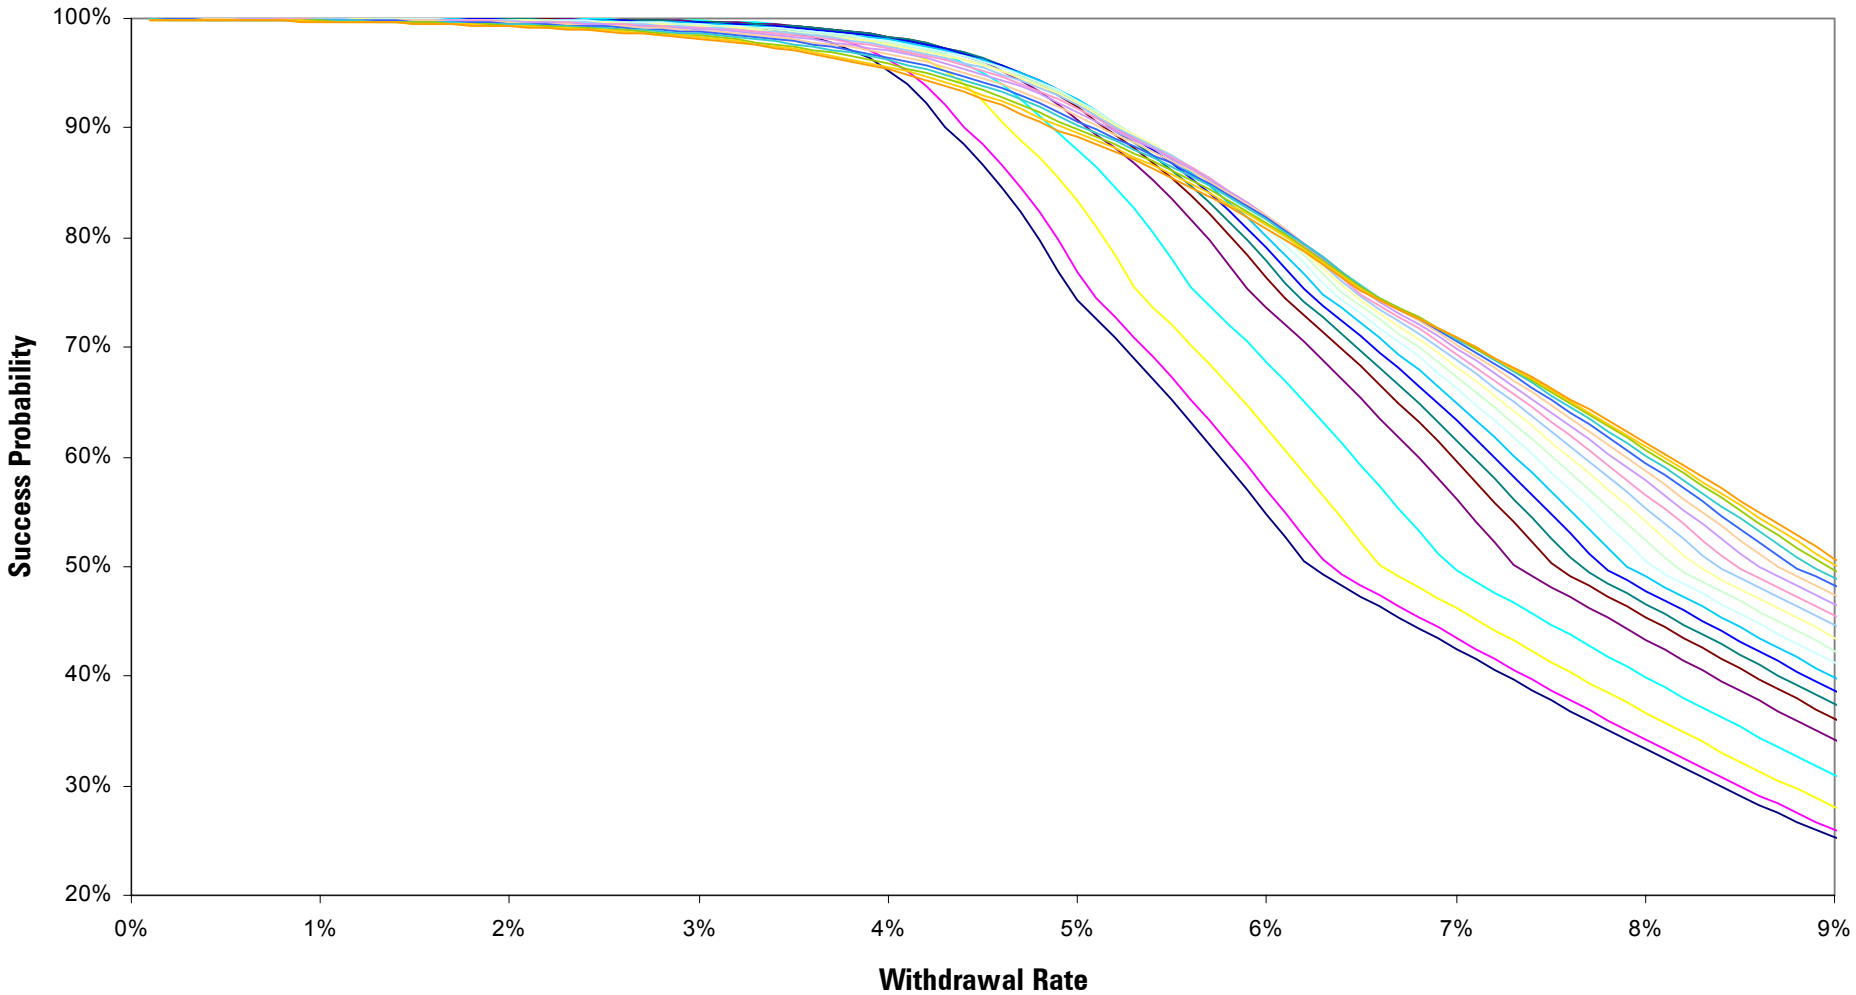# Figure 5: Withdrawal Rate vs. Success Probability at 20-Year Horizon
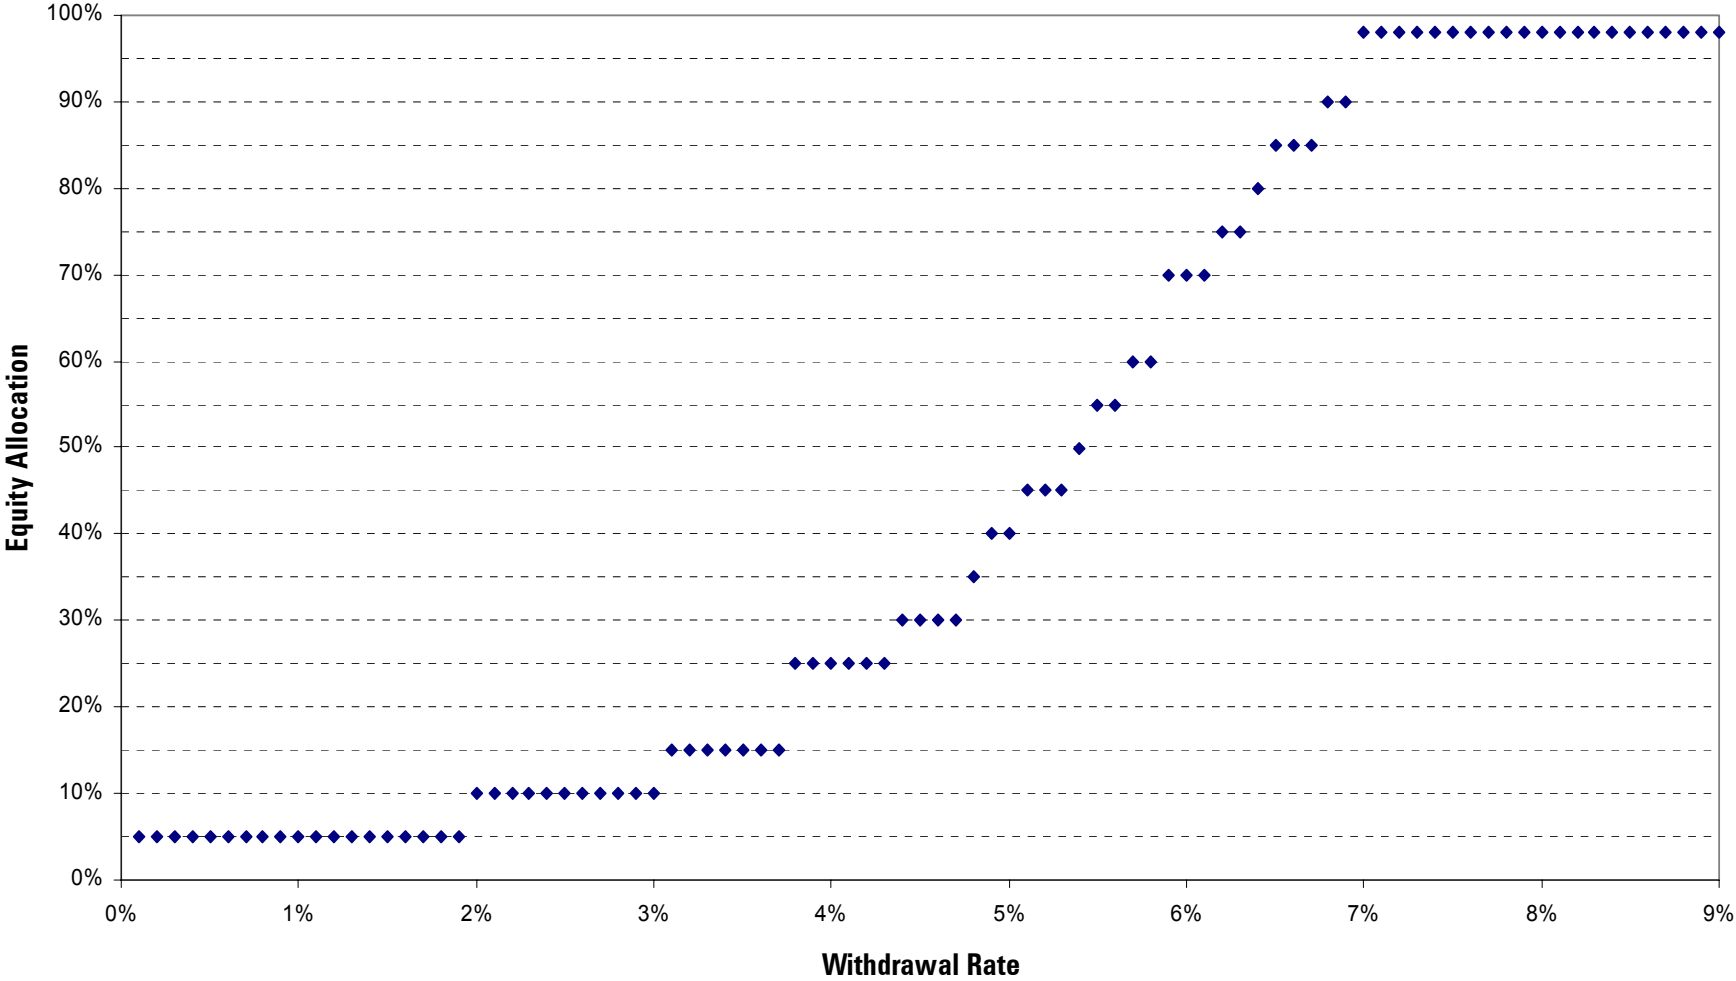
Figure 6: Mortality Rates

Figure 7: Death Probabilities for a 65-Year Old Male

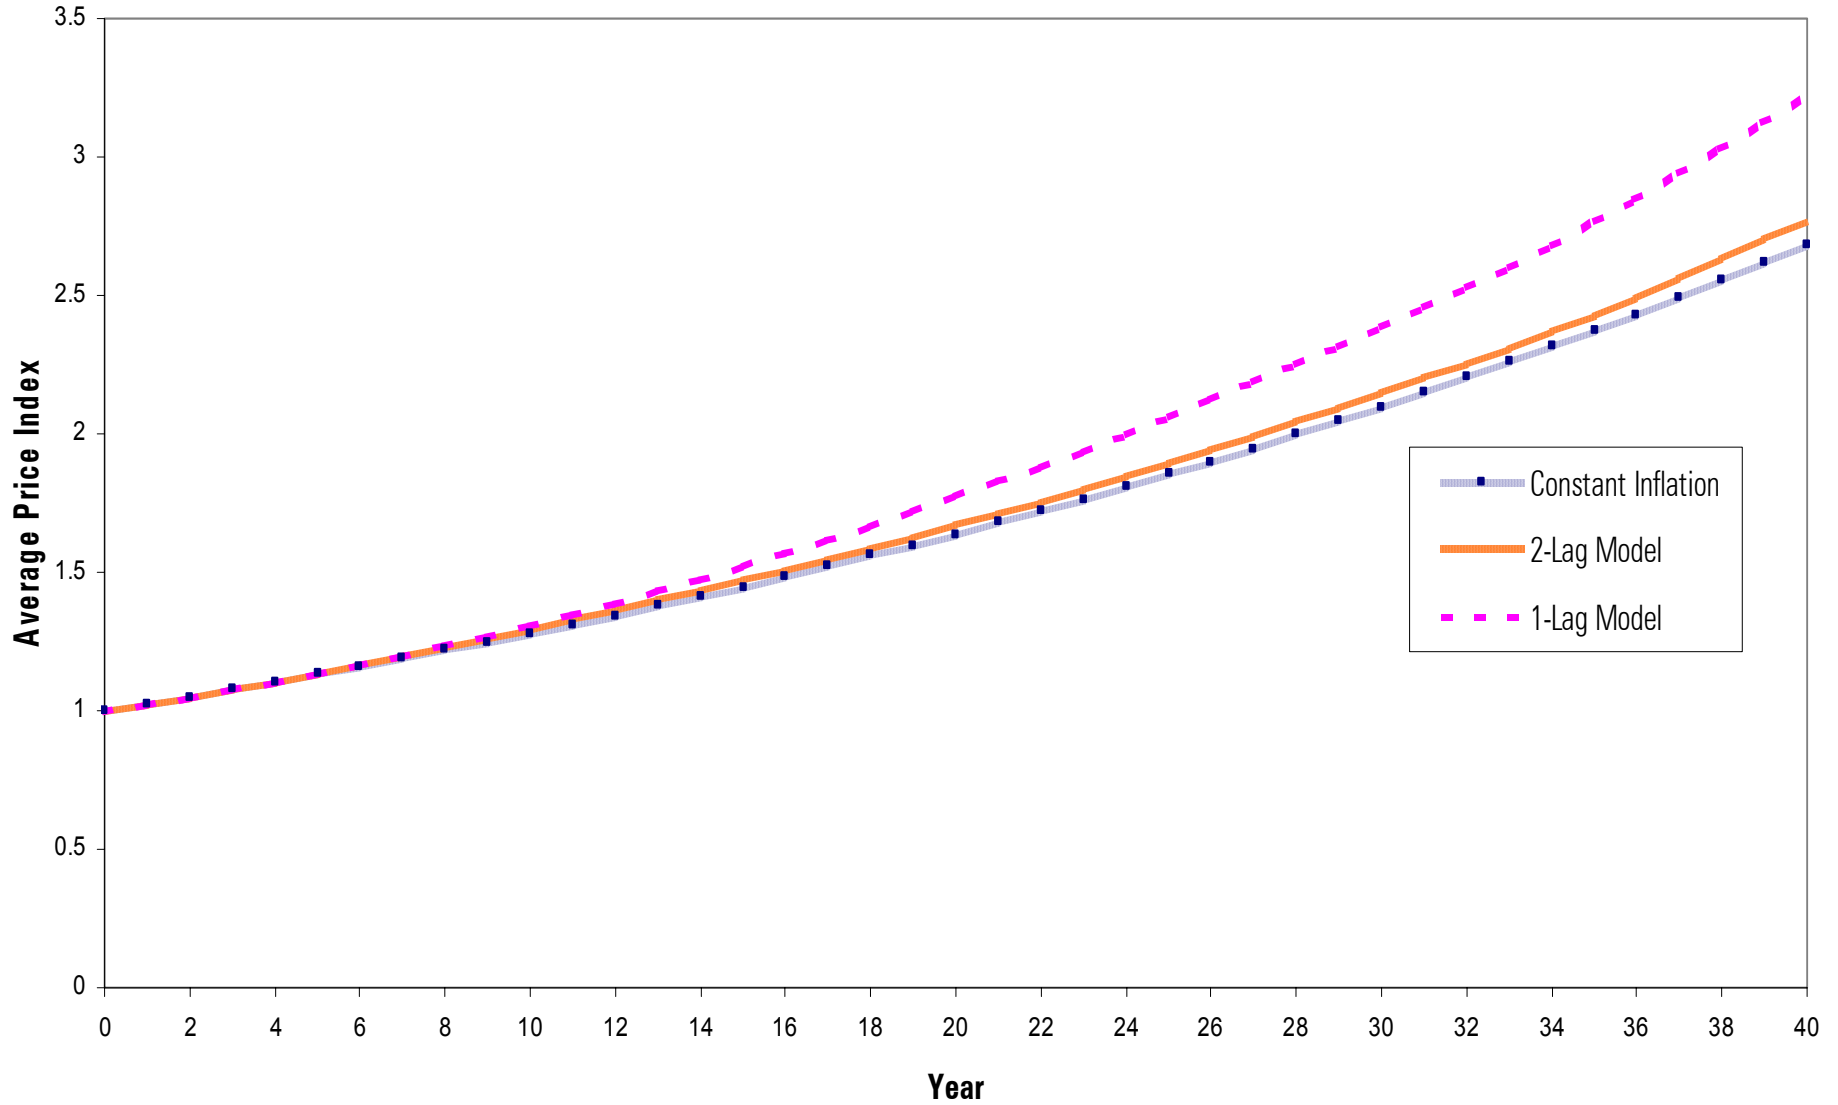**Figure 8: Withdrawal Rate vs. Success Probability: 21 Asset Mixes**

Figure 9: Equity Allocations of Mixes with Highest Success Probabilities

Equity Allocation (y-axis, 0% to 100%)

Withdrawal Rate (x-axis, 0% to 9%)

Figure 10: Average Price Indexes with Alternative Inflation Models

Figure 11: Autocorrelation Functions of Alternative Inflation Models

Number of Years
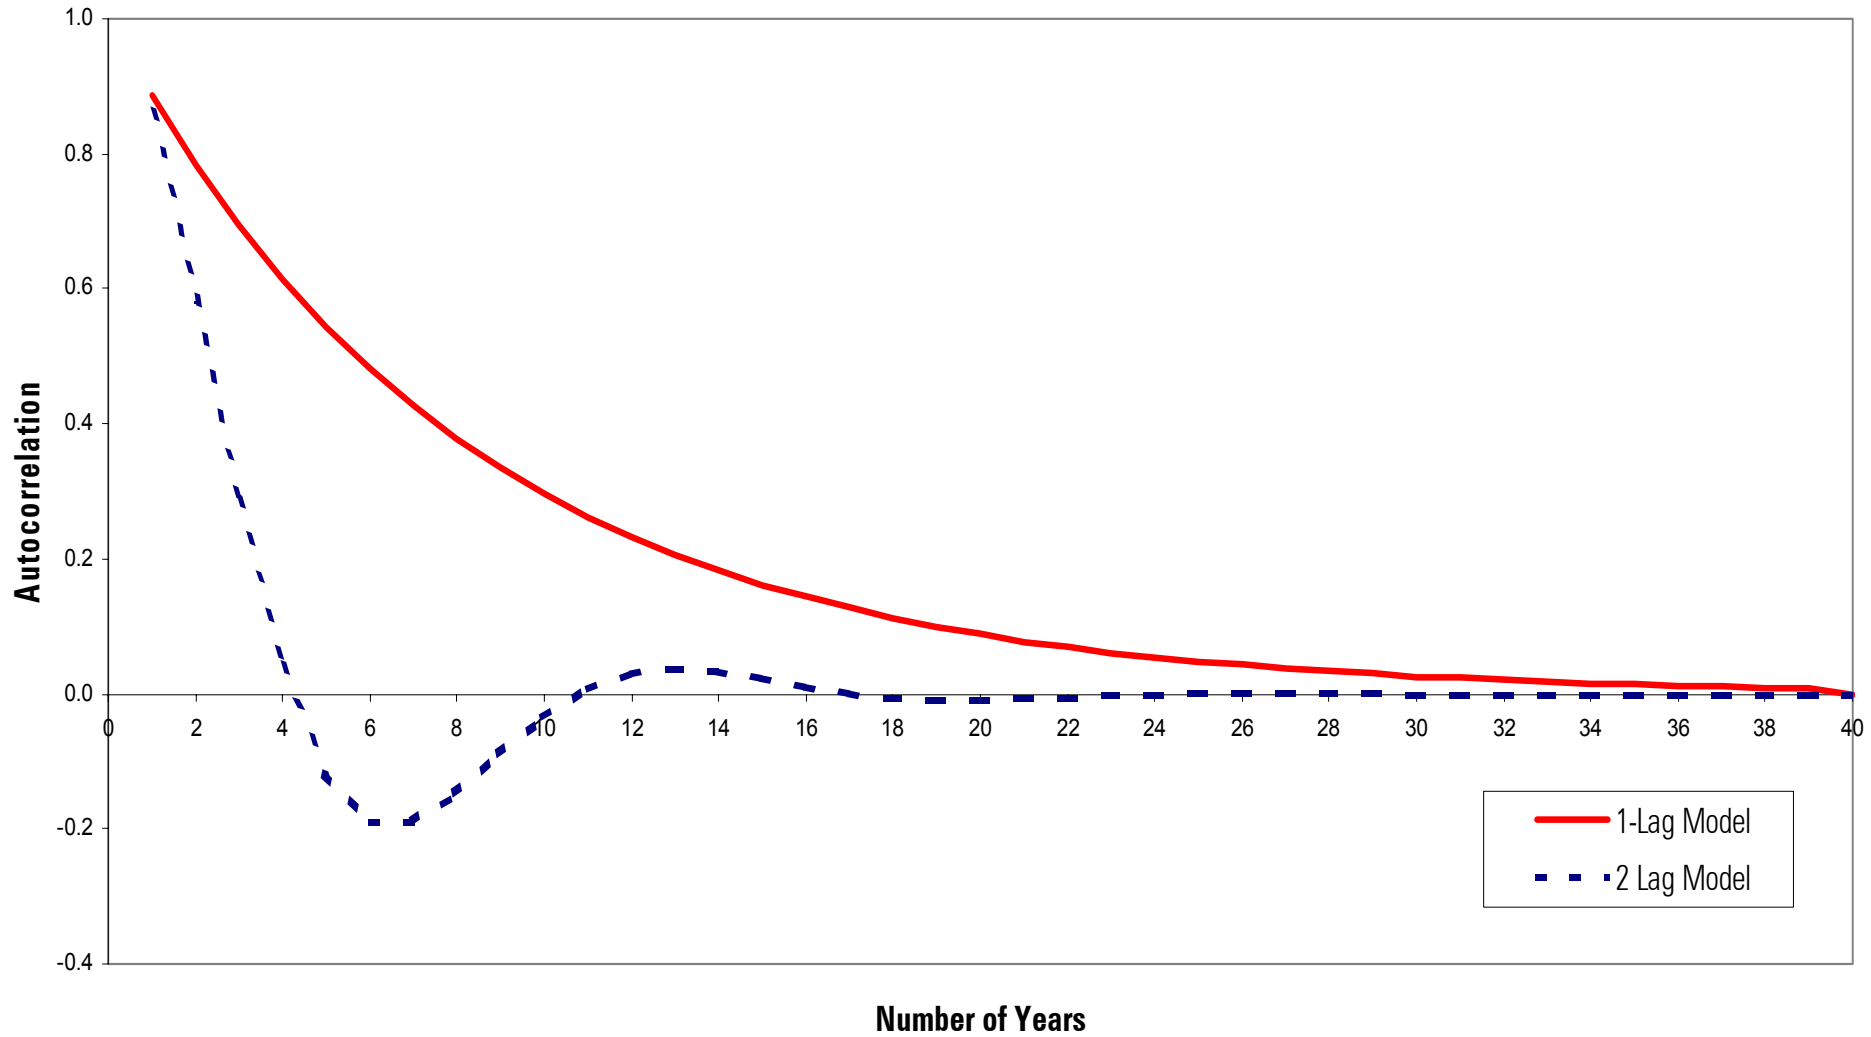
Figure 12: Forecast Volatilities for Alternative Inflation Models
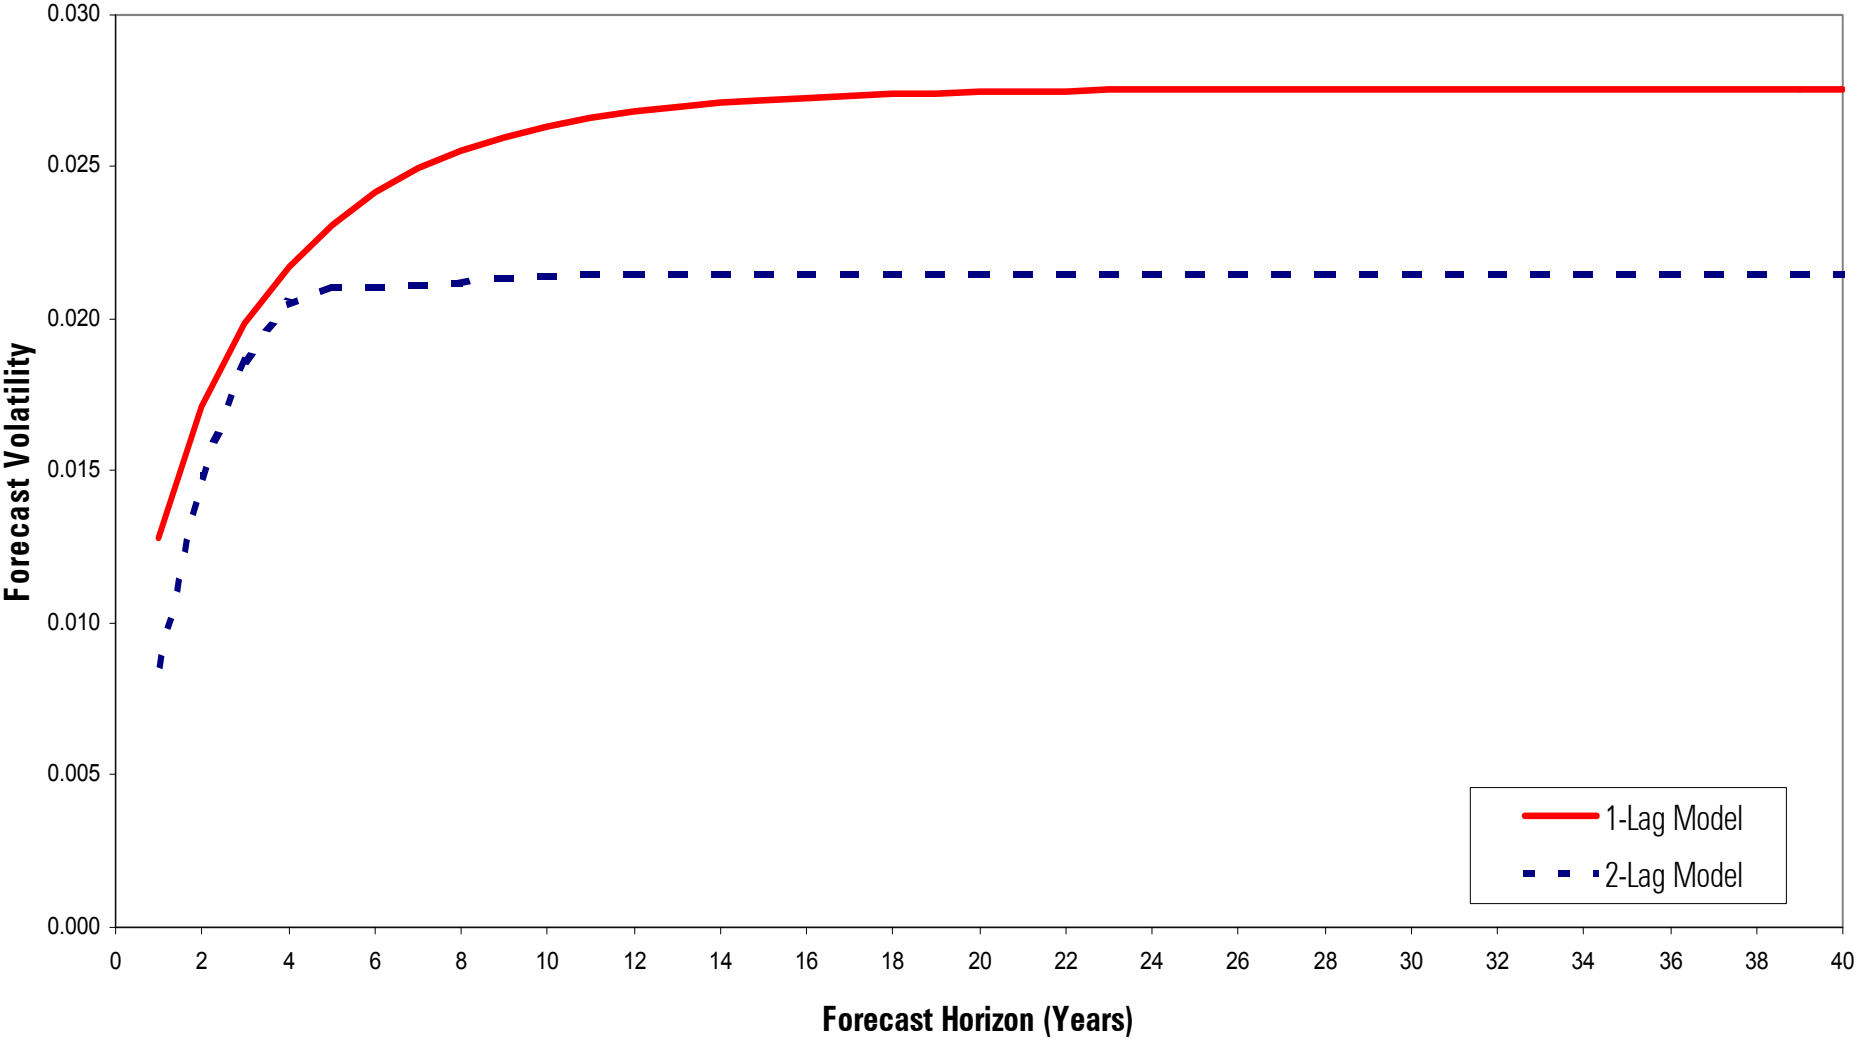
Figure 13: Success Probabilities for 5% Income Level*

Highest Probability = 94.7%
70% Annuities
70% of remaining balance in equities

Highest Probability without Annuities = 92.6%
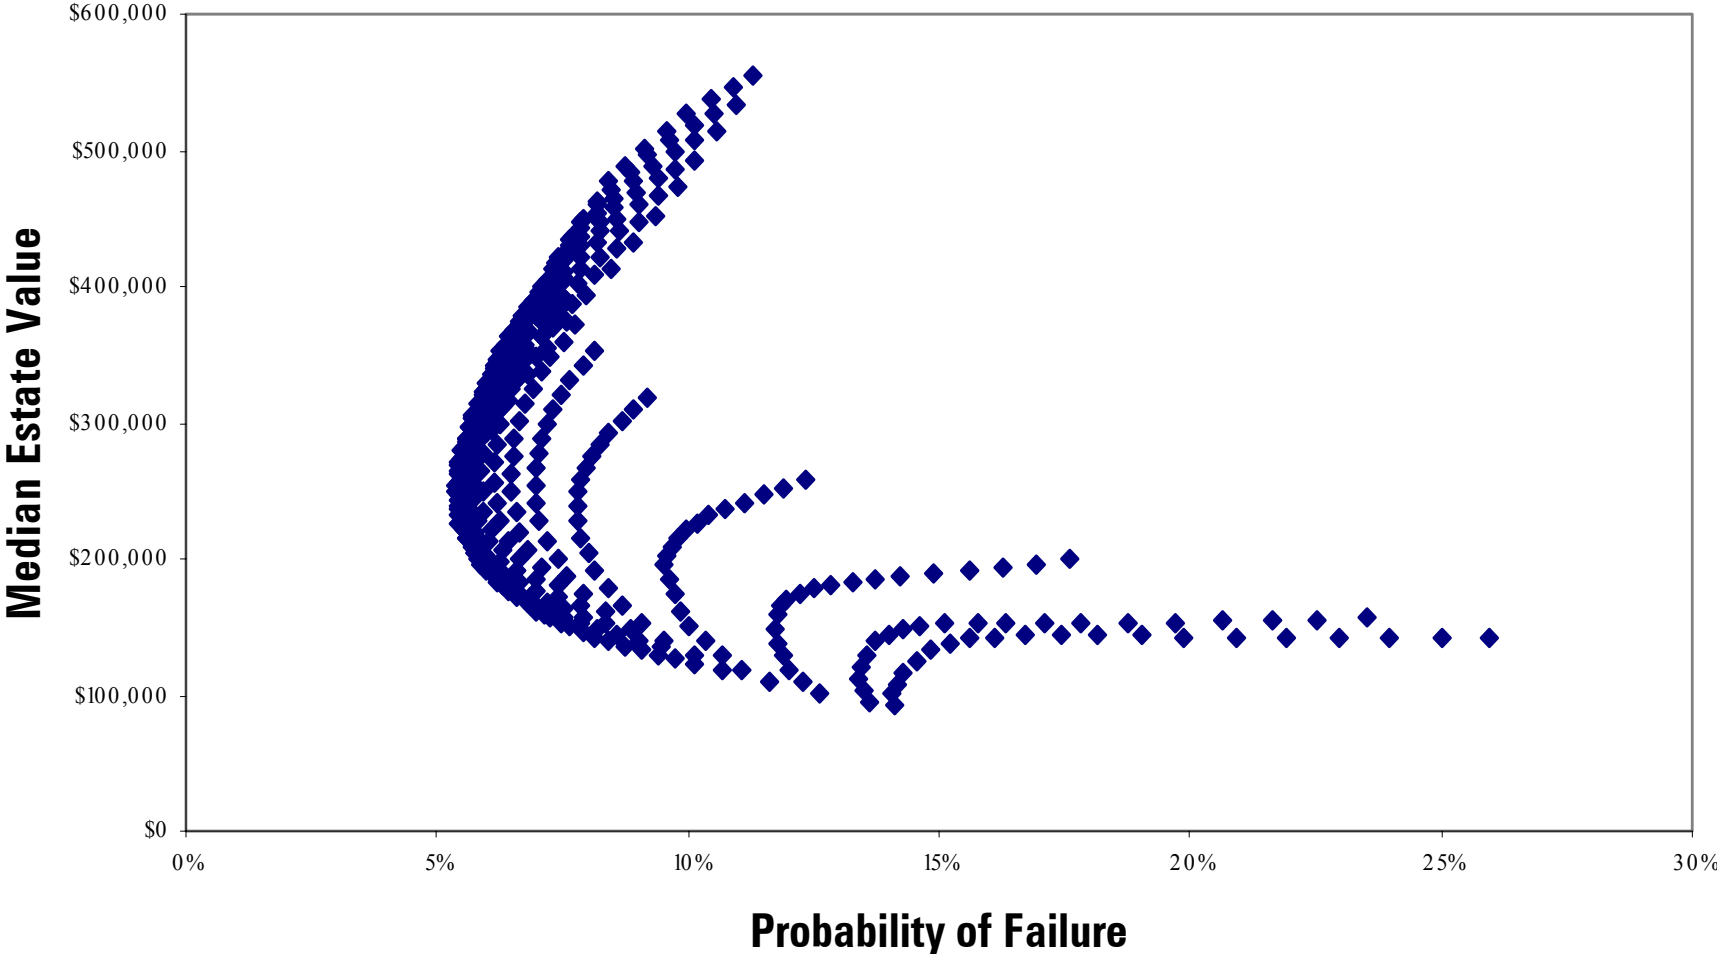40% of balance in equities

Annuity Purchase

Equity Allocation of Remaining Balance
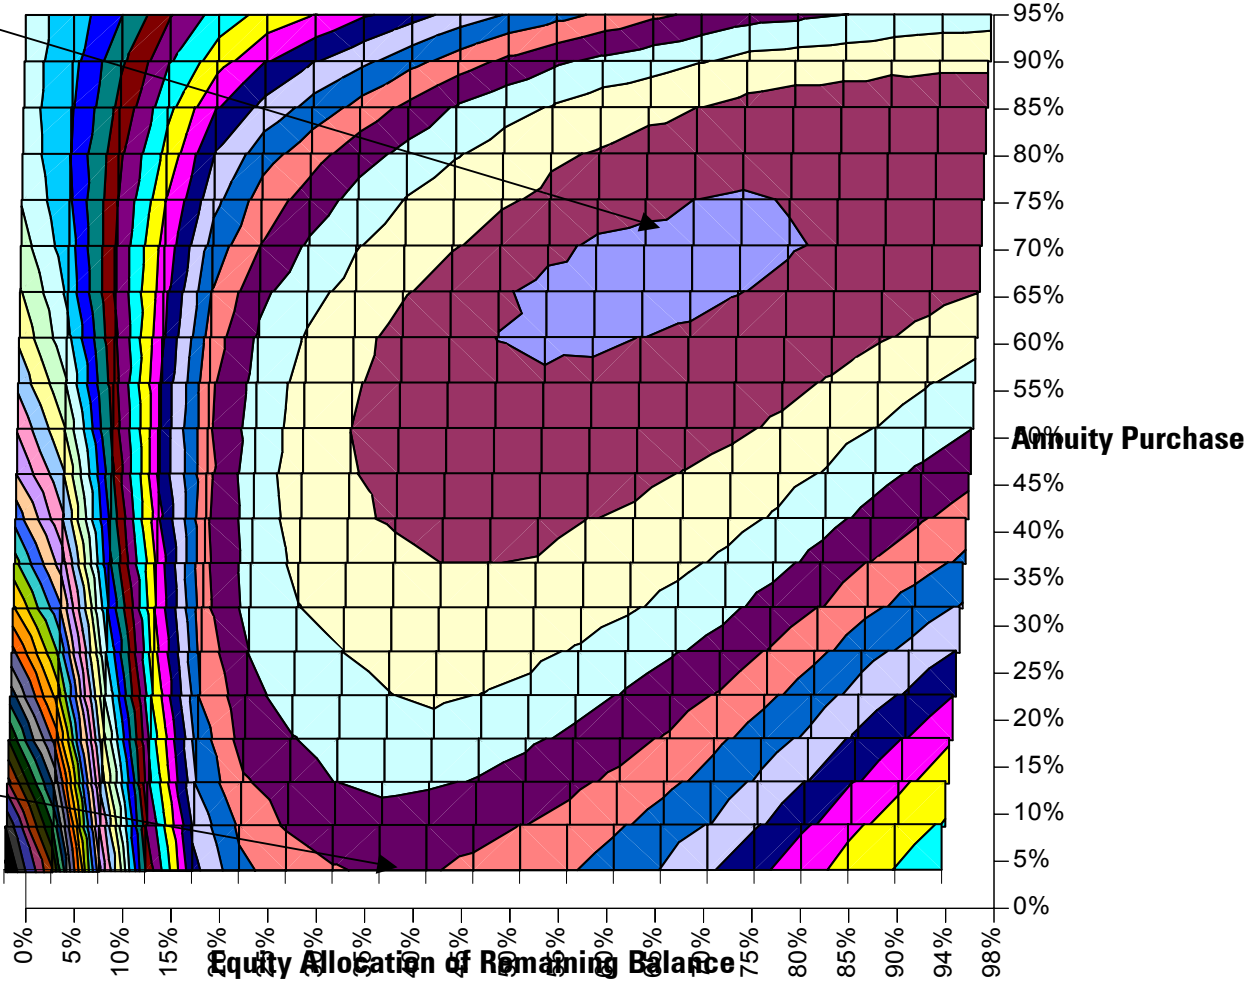
*Inflation generated by the one-lag model.

Figure 14: Highest Success Probabilities with & without Annuities, & with Alternative Inflation Models
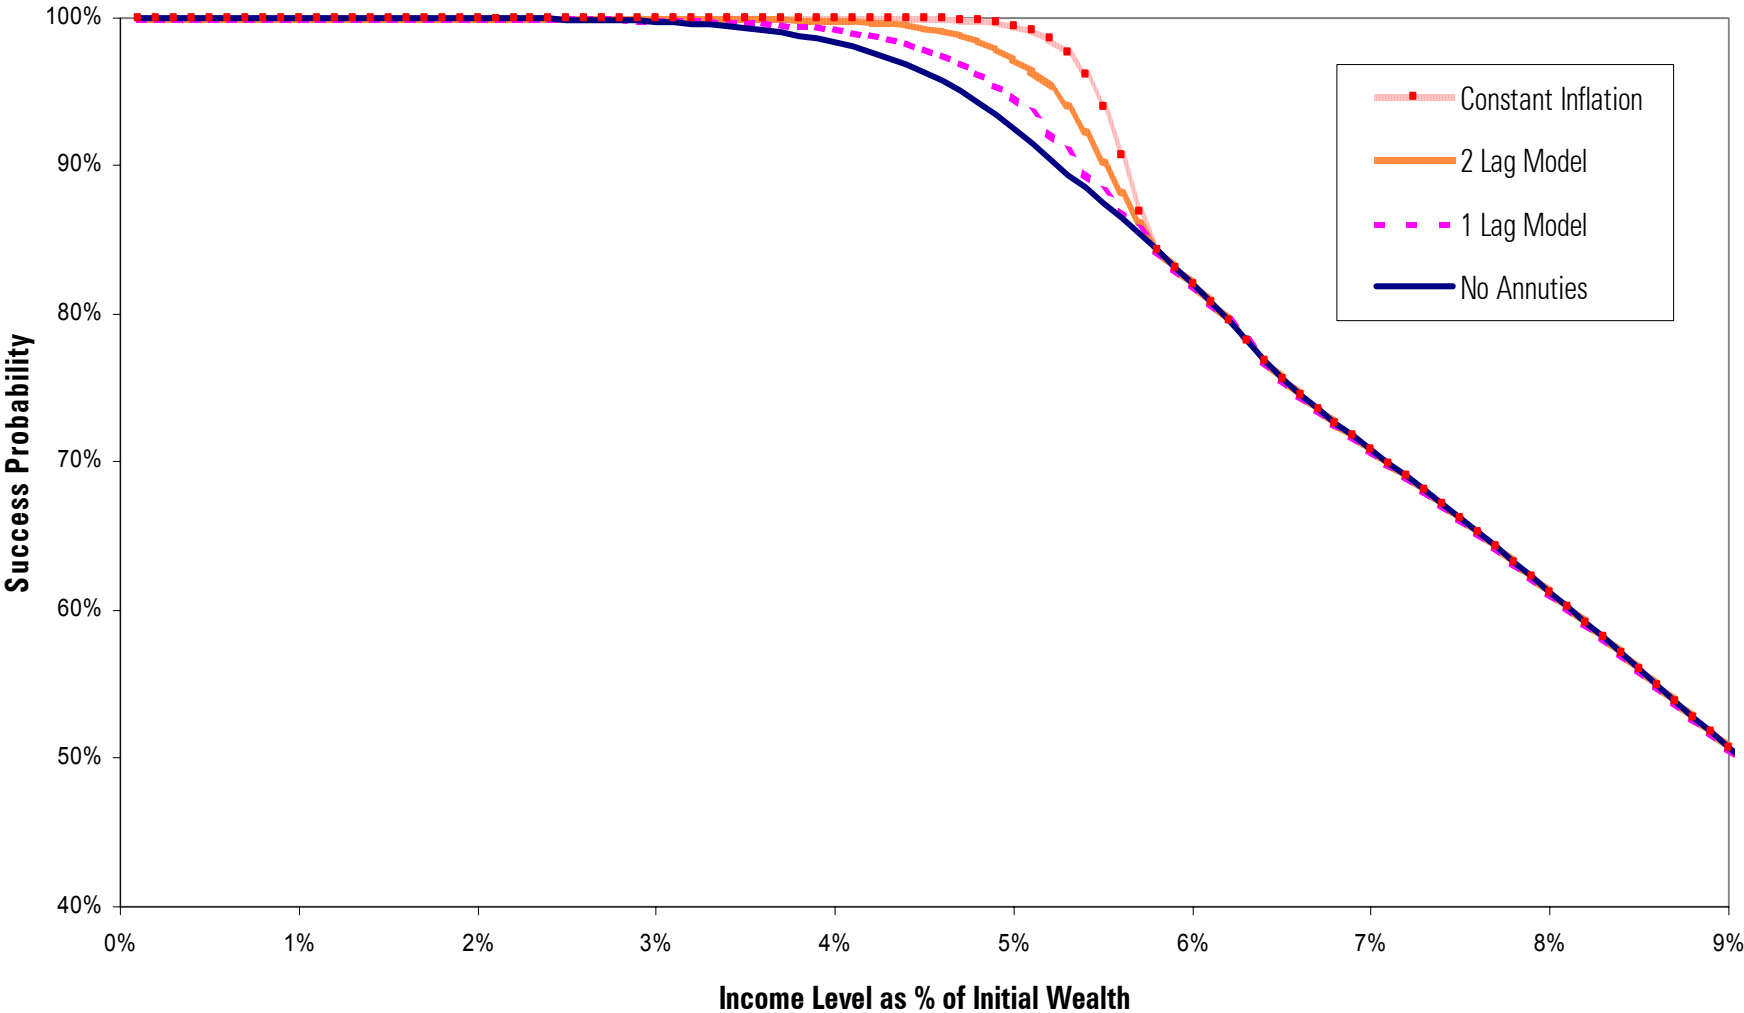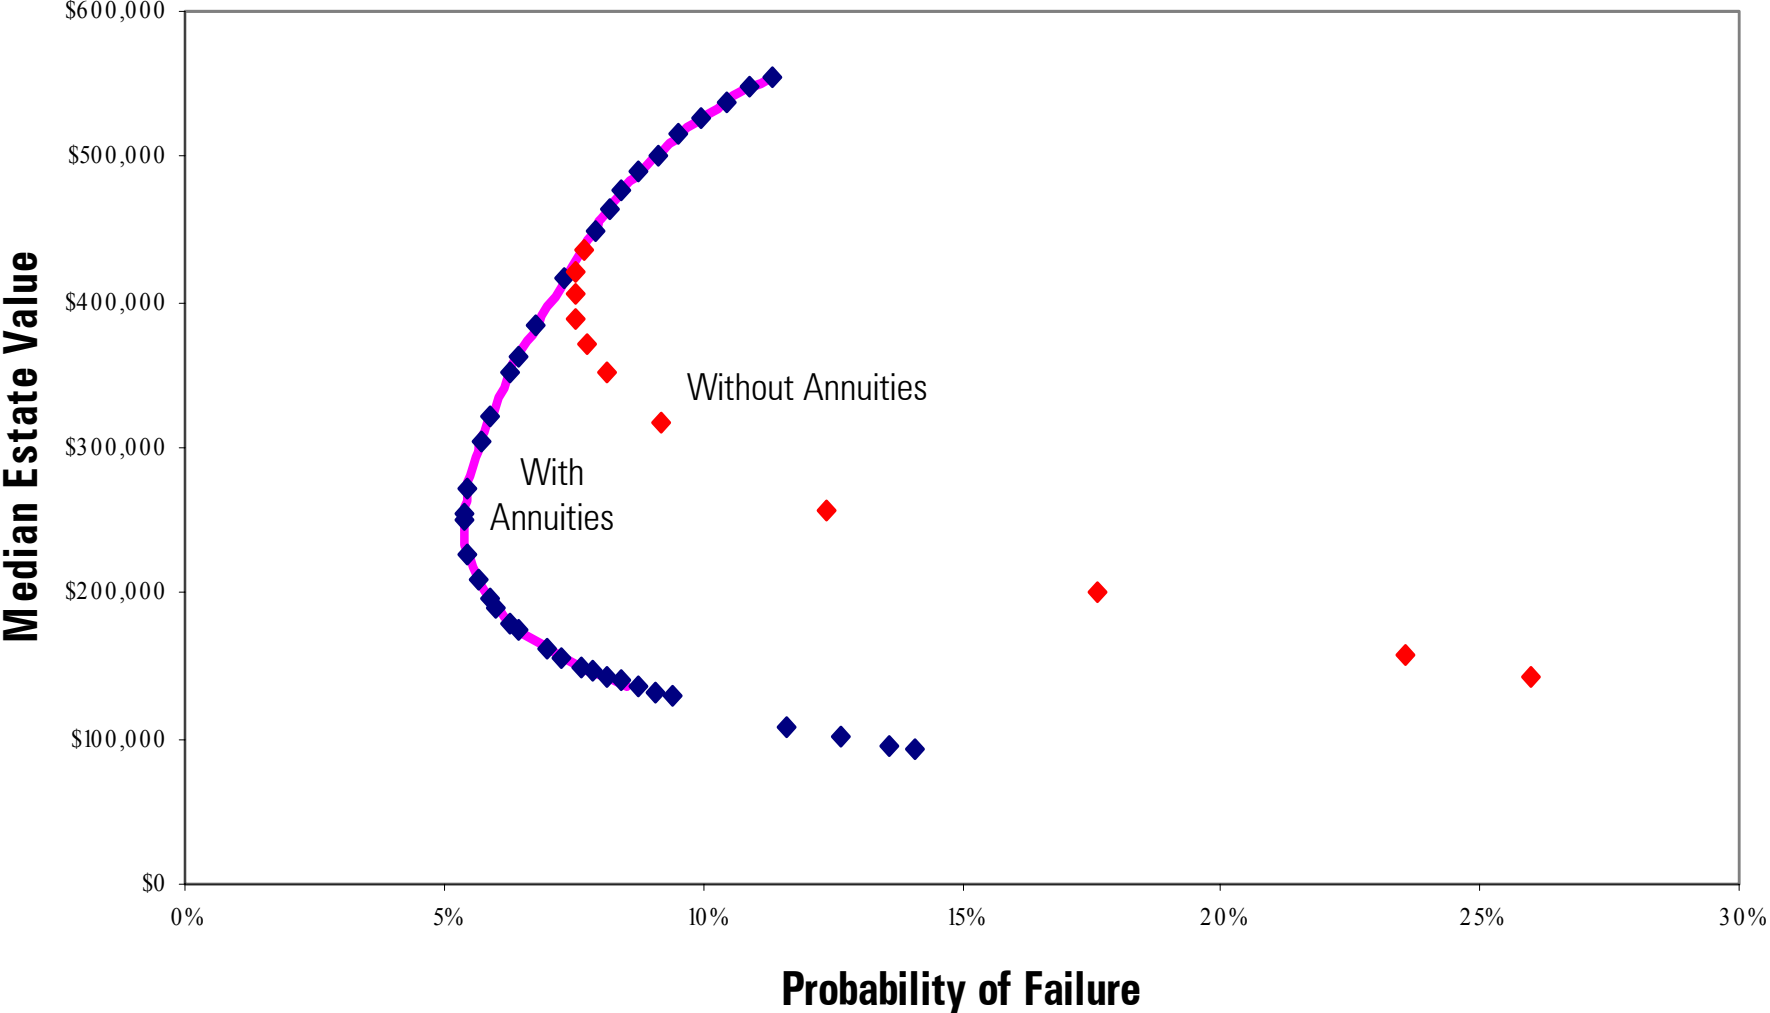
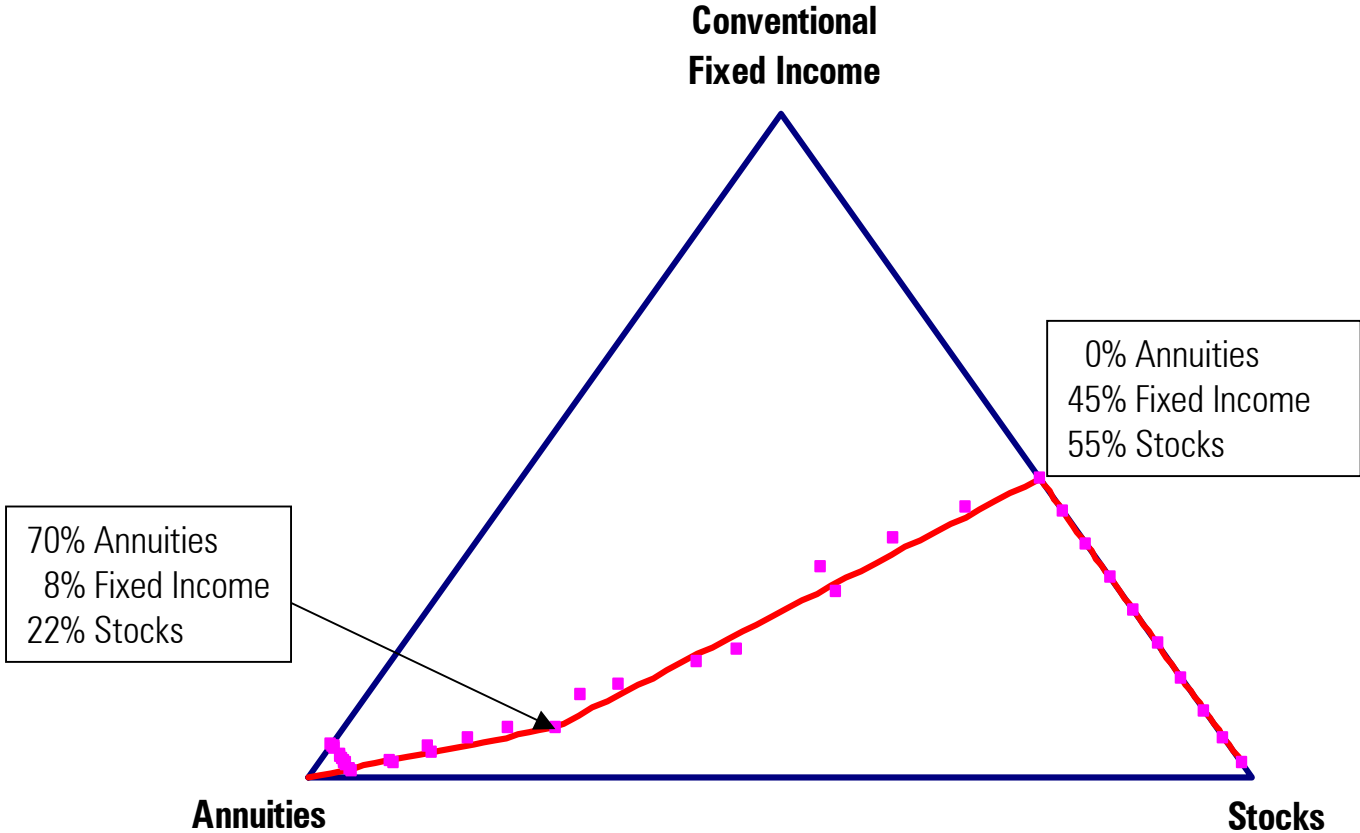**Figure 15: Failure Probability vs. Median Estate Value (5% Income Level & One-Lag Inflation Model)**

Figure 16: Failure Probability vs. Median Estate
Efficient Frontier: With & Without Annuities

Figure 17: Effective Asset Class / Annuity Combinations

Conventional
Fixed Income

0% Annuities
45% Fixed Income
55% Stocks

70% Annuities
8% Fixed Income
22% Stocks

Annuities

Stocks

Figure 18: Effective Asset Class / Annuity Mixes