10.2478/fprj-2024-0002

COMPREHENSIVE FINANCIAL PLANNING FOR A SUSTAINABLE RETIREMENT†

Nabil Tahani[1] and Chris Robinson*[2]
[1]Associate Professor of Finance and Area Coordinator, School of Administrative Studies, York University
[2]Professor Emeritus of Finance and Senior Scholar, School of Administrative Studies, York University

*Corresponding author:
Chris Robinson
Professor Emeritus of Finance and Senior Scholar, School of Administrative Studies, York University
Email: crobinso@yorku.ca

Key words:
comprehensive financial planning, sustainable retirement, probability of ruin; stochastic present value; stochastic future value

JEL Classification:
D14, D91, G11, C63

ABSTRACT

We demonstrate a comprehensive approach to the dual problems of saving enough for retirement and determining how much saving is enough at retirement date, using a probabilistic (stochastic) model. Our model estimates the probability of success of a comprehensive plan that spans the life cycle. Retirement plans are riskier than the financial planners' deterministic models show, using a simple case study as a comparative example. Variables the family can control for a successful retirement plan are the patterns of pre-retirement saving and retirement consumption, the probability of running out of money in retirement that the family will accept, and the age of retirement.

†We presented an earlier version of this paper at PFIS in Brisbane, 2020.

1. Introduction

We set the table for our model with behavioural observation. In a South African context, Willows (2020) finds that people who develop a retirement plan and stick to it are more likely to be satisfied with their financial situation and will perform better financially. Or, to use an old aphorism from corporate and personal planning, 'If you fail to plan, you plan to fail.'

Let us start with a simple case study to show how our model will provide useful information to make a successful retirement planning decision. A single man who is 30 years old meets with a financial planner to develop a savings plan to provide for retirement income. All the estimates are in real dollars. He plans to retire at age 65. He works with the planner and uses his current spending to estimate that he can save $15,000 p.a., given his expected advancement in his current occupation, the savings number could increase by 2% p.a. At age 65, he will receive $30,000 in indexed pensions, but he wants a pre-tax income of $80,000 p.a., so his savings must provide a draw of $50,000 a year when he retires. He understands enough about the risk to know that he may need higher-risk, higher-return investments in equity to meet his retirement goal.

The planner compounds the planned savings to age 65, using the future value of a constant growth annuity with a return of 5.62% p.a. and growth of 2% p.a., all in real terms. The asset allocation is approximately 90% equity and 10% risk-free assets. The underlying values come from the history of various asset classes that all finance textbooks and planners use (Ho et al. 2022, Appendix D). This produces an expected wealth at age 65 of $1.98 million. The planner discounts the $50,000 required retirement income in six different scenarios for periods of 25, 30, and 35 years (i.e., to ages 90, 95, and 100) and rates of return of 5.62% p.a. and 3% p.a. This process yields the present value of the future consumption that the accumulated savings must finance. The results range from $0.7 million to $1.1 million, which are far less than the estimated wealth accumulated to finance retirement. In addition, the planner has read Bengen (1994), who finds that 4% of the initial wealth drawn down in the first year is likely sustainable, and 4% of $1.98 million is about $79,000 p.a., far above the desired $50,000. The planner concludes this is an excellent plan with little risk involved.

This procedure fails to consider risk effectively, given the significant variation in returns, spending, and savings before retirement, in consumption during the retirement years, and the uncertain date of death. Our paper provides more information by making the critical variables stochastic (in other words, probabilistic) in a comprehensive model of retirement planning covering the entire life cycle. The model estimates the probability that a comprehensive lifetime plan will succeed. It incorporates two existing models into a single model, which is the major innovation of the paper. We are not aware of any other work that does this. The reader should understand that the deterministic model we showed above, our stochastic model, and others do not give answers or directions to act. They provide information about the risks and returns of a financial plan that consumers and their planners can use to help them decide what to do.

2. Relevant Research

We find little work on the comprehensive lifetime retirement problem that provides specific information a planner could use, starting during the saving or accumulation years and running to the end of the retirement years, during which the retiree spends the accumulated wealth.

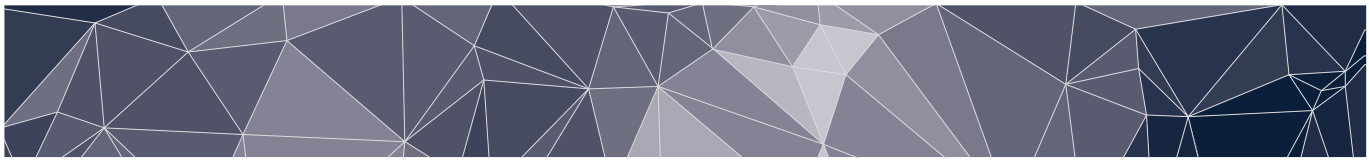

Yaari (1965) uses a general maximization of the utility of the lifetime consumption model and relates it to lifetime annuities, with an uncertain date of death but no other uncertainty. Lachance (2012) picks up on that model and imagines a realistic consumption pattern that shows it is optimal to defer retirement savings during the early years of the life cycle.

2.1 Sustainable Retirement Income

Most of the research on sustainable retirement income, also called the decumulation problem, aims to find the probability of ruin – the probability that the money runs out before the retiree(s) die. Bengen (1994) has influenced practice more than any other author, finding that 4% p.a. is a sustainable withdrawal rate for a person retiring at age 65, using random draws of historic return patterns. Bajtelsmit & Wang (2018) maximise the utility of consumption to focus on longevity risk. They posit a case study with specified death dates and use simulation to study the effect of various practical responses to longevity risk like annuities, deferred retirement, and reverse mortgages. Their conclusion emphasizes the importance of longevity risk and is consistent with our results. MacDonald et al. (2018) investigate the effect of different risks by simulating mortality, rate of return, and inflation risk in a single case study. They find that deterministic financial plans are substantially riskier than they appear.

Several papers use simulation of returns with varying assumptions about time horizon, some of which we cite. Drew & Walk (2015) find the 4% rule is unlikely to be sustainable using return history from five countries and simulation. Their result is significant in their Australian context, where most retirement income falls under the Superannuation scheme, which does not impose annuitisation of the accumulated balance at retirement. In countries like Canada and the United States (US), government and employer pension plans provide a life annuity at retirement. Suarez et al. (2015) and Suarez (2020) use simulation of rate of return and historical values to investigate failure, including sequencing risk. Finke et al. (2013) find that target withdrawal rates previously specified in other simulation studies are not sustainable in the return environment current at that time, but fixed income rates have risen considerably since 2013.

Estrada & Kritzman (2019) approach the problem somewhat differently from all the other literature to provide more information for the decision process. They create a measure they call the coverage ratio. In their own words (Estrada & Kritzman, 2019, pg. 36):

> Formally, let Y_t be the number of years of inflation-adjusted withdrawals sustained by a strategy, during and after the retirement period, and L be the length of the retirement period considered. Then we define the coverage ratio in retirement period t (C_t) as $C_t = Y_t / L$.

A value of $C < 1$ is a failure and $C > 1$ implies a future bequest. Estrada (2023) expands on this to look at deciles of coverage ratios with also deciles of % of equity invested (e.g. 80% equity, 20% bonds) using rolling 30-year historic returns for many different countries.

All the studies discussed so far concentrate on making the rate of return a random variable in one way or another and then comparing the results for different lengths of retirement period. Robinson & Tahani (2010) provide an analytic solution to the problem of sustainable retirement income that incorporates stochastic date of death, stochastic rate of return, correlation between

the returns and the consumption, and the drift and volatility of the amount of consumption. The planner specifies the real return distribution and the consumption parameters based on the client's desired consumption patterns in retirement. The model produces a *stochastic present value* (SPV), which is not a number but a distribution of a random variable and shows the probability that the withdrawal rate is sustainable. This paper extends Milevsky & Robinson (2005) by making consumption a stochastic variable instead of a fixed value.

2.2 Pre-retirement Savings Optimisation

Tahani & Robinson (2010) is the only paper to develop and investigate the properties of a model of the pre-retirement savings optimisation problem using a stochastic model with more than just the rate of return risk included. The client specifies the retirement goal and date as well as the initial periodic savings. The initial value of savings is already known. The planner determines the mean and standard deviation of a lognormal rate of return, with the parameters suitable to a given risky asset allocation (e.g., a balanced portfolio, all equity, etc.) and a risk-free rate. The planner estimates the drift rate, the volatility of the periodic savings amount, as well as the correlation between the rate of savings and the rate of return. The model produces a *stochastic future value* (SFV), which is not a number but a distribution of a random variable, and estimates the probability that the goal will be met for the periodic savings amount and any set of values for the other variables. It also determines the allocation between the risky portfolio and a riskless asset that will minimize the probability of shortfall. The creation of the SFV improves upon earlier work by Ho et al. (1996).

3. Stochastic Modelling of Retirement Planning

The model presented in this paper involves some complex mathematics. It requires the following inputs:

- Drift means the average amount a random variable can increase/decrease (positive/negative) each year: (μ, η) are the drifts for the risky asset V_t and the annual saving $S_{t,}$ respectively.
- Volatility or diffusion is the term in stochastic calculus for what would be the standard deviation of a random variable in discrete form: (σ, β) are the diffusion terms for the risky asset V_t and for the annual saving $S_{t,}$ respectively.
- Initial wealth that will accumulate until retirement: w.
- Initial annual savings amount, a random variable with a drift and a volatility the planner or client estimates: S_t.
- Number of years from now to a planned retirement date and age at planned retirement, which the client determines: T.
- Median age of death for someone who retires at the planned date, based on the mortality rate λ. This value depends on the country or could even be specific to a client who knows more about her health relative to the average.
- The acceptable probability of running out of money during retirement. The client sets this number according to her comfort level. It cannot be zero, though in truth, we suppose that for some huge fortunes, it is effectively zero.

- The real dollar amount of annual consumption, C_t, desired in retirement with drift and diffusion (κ,ψ).
- The cumulative return is based on a portfolio consisting of a fraction of $(1-\alpha)$ invested in the risk-free asset and a fraction α invested in the risky asset (V_t).
- A correlation between savings amounts each year and investment return, ρ; and a correlation between consumption in retirement each year and investment return, ρ'. Although savings amount and retirement income are primarily related to other factors, in general, we expect a person to be able to save more if investment returns are high and to consume more in retirement if the retirement portfolio performs better than expected, and conversely, if investment returns are below expectation.

What does the model produce?

- It determines the optimal allocation, $\hat{\alpha}$, between the risky and riskless asset subject to the specified constraint of the acceptable probability of running out of money during retirement. This allocation is constant, though the realised return is still a random variable.
- It produces the probability that the entire plan from today through the entire retirement period will succeed – the retiree will not run out of money.
- It produces the expected level of wealth at retirement date, which will not be identical to the expected wealth calculated using the deterministic model that we demonstrated in the Introduction. However, it will not be too different.

Figure 1 below summarizes the comprehensive model's inputs and outputs. Although it might look like there are several steps, everything is optimised at once.

Figure 1. Inputs and Outputs for the Model.

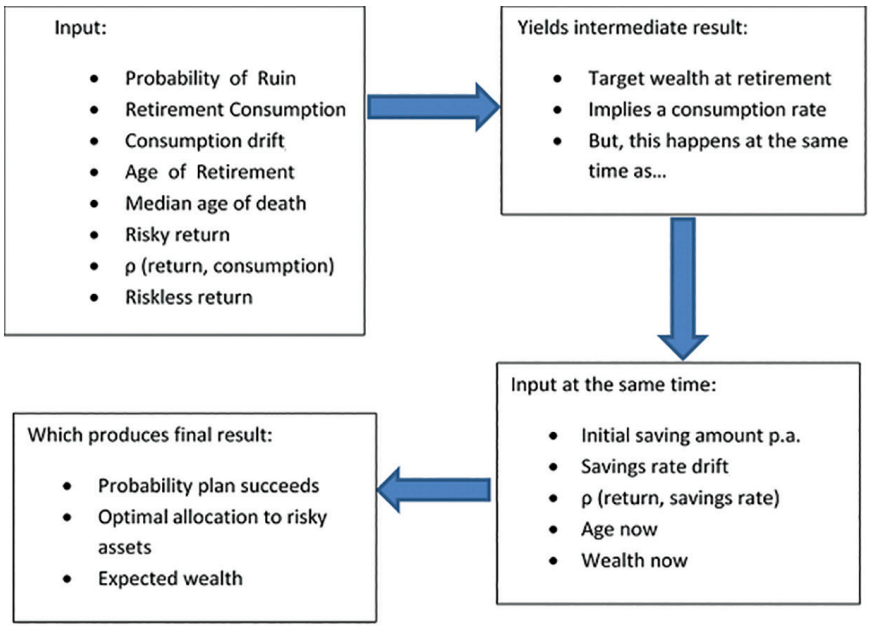

We revisit retirement plans periodically and revise them as circumstances change in practice. Our model provides a fast look at how the plan is progressing towards the targets. Every time new information is added, the optimal asset allocation changes, providing guidance if rebalancing is needed.

3.1 Pre-Retirement Model: Cumulative Return, Initial Wealth and Stochastic Savings

By analogy to the stochastic present value in Milevsky and Robinson (2005) and Robinson and Tahani (2010), we introduce the notion of the *stochastic future value* (SFV) of the stochastic savings and initial wealth at a known future date (e.g., retirement date). The SFV is a random variable that can only be known through its probability distribution. More specifically, we define it by (technical details are presented in Appendix A):

$$SFV = w\,R_T + R_T \int_0^T S_t R_t^{-1} dt \tag{1}$$

where T denotes the known retirement date (in years), w is the initial wealth, and (S_t) and (R_t) are the savings and cumulative return processes. We assume that the processes (S_t) and (R_t) follow two correlated Geometric Brownian motions given by:

$$\begin{cases} dV_t = \mu V_t dt + \sigma V_t dB_t & ; \quad V_0 = v \\ dR_t = \tilde{\mu} R_t dt + \tilde{\sigma} R_t dB_t & ; \quad R_0 = 1 \\ dS_t = \eta S_t dt + \beta S_t dZ_t & ; \quad S_0 = s \\ d < B, Z >_t = \rho\, dt \end{cases} \tag{2}$$

where (V_t) is the risky asset process. The constants (μ, η) and (σ, β) denote respectively the drifts and the diffusion parameters of (V_t) and (S_t). The last equation states that the correlation between the savings and the cumulative return (or the risky asset) is equal to ρ. Since the cumulative return is based on a portfolio consisting of a fraction of $(1-\alpha)$ invested in the risk-free asset and a fraction α invested in the risky asset (V), we have:

$$\begin{cases} \tilde{\mu} = (1-\alpha)\, r + \alpha\mu \\ \tilde{\sigma} = \alpha\sigma \end{cases} \tag{3}$$

where r is the risk-free rate of return. We need the optimal value $\hat{\alpha}$ that maximizes the probability of reaching the retirement goal G, which is $\hat{\alpha} = \underset{\alpha}{argmax}\, P\left(SFV(\alpha) \geq G\right)$. Using standard stochastic calculus, it can be shown that the solutions to Eq. (2) are given by:

$$\begin{cases} R_t = R_0 \exp\left(\left(\tilde{\mu} - \frac{1}{2}\tilde{\sigma}^2\right)t + \tilde{\sigma} B_t\right) \\ S_t = S_0 \exp\left(\left(\eta - \frac{1}{2}\beta^2\right)t + \beta Z_t\right) \end{cases} \tag{4}$$

To value the probability of reaching the retirement goal, we need the probability distribution of the SFV as given in Eq. (1). Unfortunately, the SFV is similar to a (continuous) sum

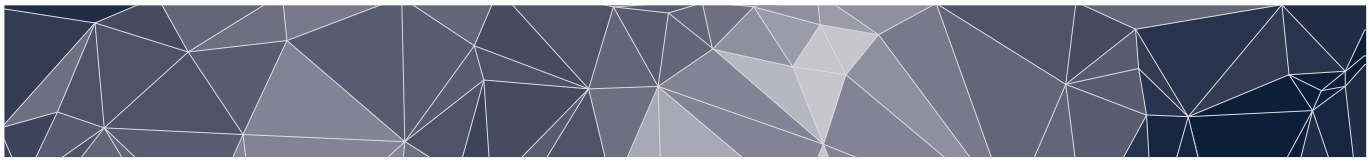

of lognormal variables for which there is no simple closed-form cumulative distribution function. Thanks to the Asian options (i.e., arithmetic average options) literature, many analytical approximations based on known probability distributions are available. We will approximate the SFV distribution using the lognormal distribution by matching the first two moments.

Main result: Lognormal approximation

Based on the first two moments, E_1 and E_2 of the SFV (presented in Appendix A), we can approximate the probability of reaching the retirement goal by:

$$P\left(SFV \geq G\right) = 1 - N\left(\frac{\ln(G) - a}{b}\right) \tag{5}$$

where $N(.)$ is the standard normal cumulative distribution and the parameters a and b are given by:

$$\begin{cases} a = 2\ln(E_1) - \frac{1}{2}\ln(E_2) \\ b = \sqrt{\ln(E_2) - 2\ln(E_1)} \end{cases} \tag{6}$$

Knowing the properties of the standard normal distribution, we can calculate the probability of reaching the retirement goal in Eq. (5) as a function of the fraction α, and then maximize it to determine the optimal allocation $\hat{\alpha}$, the corresponding probability of reaching G and the expected wealth at the retirement date.

3.2 Post-Retirement Model: Shortfall Probability, Retirement Consumption

We expand the Milevsky & Robinson (2000, 2005) model of retirement consumption to incorporate a stochastic consumption rate. This setting allows for more flexibility, and different consumption patterns can be modeled within the proposed framework. We create the *stochastic present value* (SPV) of the annual stream of stochastic consumption at the retirement date. The SPV is like the present value of a stream of cash flows, but it is a random variable that can be known only through its probability distribution. We define it by (technical details are presented in Appendix B):

$$SPV = \int_0^{\tilde{T}} C_t R_t^{-1} dt = \int_0^{+\infty} 1_{(\tilde{T}>t)} C_t R_t^{-1} dt \tag{7}$$

where $\tilde{T}$ denotes the random time of death (in years), (C_t) and (R_t) are respectively the consumption and the cumulative return processes (defined previously, see Eq. 2), and $1_{(\tilde{T}>t)}$ is the indicator function that equals 1 if $\tilde{T} > t$, and 0 otherwise. We assume that the process (C_t) is correlated with (R_t) and is defined by:

$$\begin{cases} dC_t = -\kappa C_t dt + \psi C_t dZ_t \quad ; \quad C_0 = 1 \\ d < B, Z >_t = \rho' dt \end{cases} \tag{8}$$

where the last equation states that the correlation between the consumption and the cumulative return is equal to ρ'. Using standard stochastic calculus, the solution to equation (8) is given by:

$$\begin{aligned} C_t &= \exp\left(-(\kappa + \frac{1}{2}\psi^2)t + \psi Z_t\right) \\ &= \exp\left(-(\kappa + \frac{1}{2}\psi^2)t + \rho'\psi B_t + \sqrt{1-\rho'^2}\psi \bar{B}_t\right) \end{aligned} \tag{9}$$

where $(\bar{B}_t)$ is a Brownian motion uncorrelated with (B_t). The last equation in Equation (9) shows how the correlation coefficient affects consumption. To value the probability of ruin using the SPV, we need to use the *discounted consumption* $(C_t R_t^{-1})$. Combining (R_t) from Equation (4) and (C_t) from Equation (9), we can show that:

$$C_t^{-1} R_t \underset{D}{\sim} \exp\left(\left(\bar{\mu} - \frac{1}{2}\bar{\sigma}^2\right)t + \bar{\sigma}\bar{Z}_t\right) \tag{10}$$

where $(\bar{Z}_t)$ is a Brownian motion and:

$$\begin{cases} \bar{\mu} = \tilde{\mu} + \kappa + \psi^2 - \rho'\tilde{\sigma}\psi \\ \bar{\sigma} = \sqrt{\tilde{\sigma}^2 + \psi^2 - 2\rho'\tilde{\sigma}\psi} \end{cases} \tag{11}$$

The inverse of the discounted consumption $\left(C_t R_t^{-1}\right)$ has the same probability distribution (and hence the same moments) as the Geometric Brownian motion with a drift of $\bar{\mu}$ and a volatility of $\bar{\sigma}$.

We want to determine how likely it is that a person will run out of money in retirement, i.e., the probability that the SPV of future consumption will exceed the wealth we have to support that consumption. To do that, we need to introduce one last stochastic variable to represent the age of death.

Exponential lifetime random variable

The remaining lifetime random variable $\tilde{T}$ is assumed to be exponentially distributed with a mortality rate of λ; that is, the probability distribution for $\tilde{T}$ is given by:

$$P(\tilde{T} > t) = \exp\left(-\lambda t\right) \tag{12}$$

The mortality rate can be fitted to a mortality table using the median value given by $Median = \frac{\ln(2)}{\lambda}$. The theoretical average (expected) lifetime is given by $Mean = \frac{1}{\lambda}$. The case $\lambda = 0$ corresponds to the endowment (infinite life) case. Although the paper focuses only on the exponential formulation of the remaining lifetime, the following result can be extended to other mortality laws such as Gompertz and Gompertz-Makeham.[1]

[1] Huang et al. (2017) note that a person's biological age may differ materially from their chronological age. Measurement of biological age is a continuing research topic and perhaps someday it will replace chronological age in annuity calculations and financial planning models.

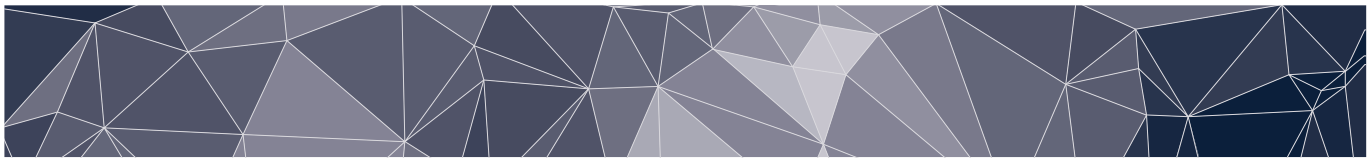

Main result: Exponential - Reciprocal Gamma approximation

Following Milevsky (1997) and Milevsky & Robinson (2005), we can show that the probability of ruin, where the SPV is higher than the wealth at retirement *G*, can be approximated with a *Reciprocal Gamma* distribution. More precisely, we have that:

$$P(SPV > G) \approx GammaDist\left(\bar{\alpha}, \bar{\beta}; \frac{1}{G}\right) \tag{13}$$

where $GammaDist\left(\bar{\alpha}, \bar{\beta}; x\right)$ is the cumulative distribution function of the Gamma distribution with parameters:

$$\begin{cases} \bar{\alpha} = \dfrac{2\bar{\mu} + 4\lambda}{\bar{\sigma}^2 + \lambda} - 1 \\ \bar{\beta} = \dfrac{\bar{\sigma}^2 + \lambda}{2} \end{cases} \tag{14}$$

For the endowment case ($\tilde{T} = +\infty$), Equation (13) gives the exact distribution with parameters $\bar{\alpha}_{\infty} = \dfrac{2\bar{\mu}}{\bar{\sigma}^2} - 1$ and $\overline{\beta_{\infty}} = \dfrac{\bar{\sigma}^2}{2}$. For the mathematical details of the proof, refer to Milevsky (1997).

We know the properties of the gamma distribution, so we can calculate the probabilities of ruin for a retirement plan in which the initial wealth and starting real consumption are given, and the mortality rate, rate of return, and pattern of consumption are all based on probabilities rather than certainties.

4. Results from Applying the Model

This model does not incorporate every aspect of retirement planning. Still, it provides a more useful evaluation of the probability of success of a comprehensive lifetime plan than any existing technique. This section explains how a planner would use this model and shows some examples. The key output is the probability the plan succeeds. A successful strategy is one in which the consumer dies before the retirement savings are exhausted. A failure does not mean starvation; it means the person or family could not consume the desired amount every year in retirement, but that shortfall might be only one dollar.

All the values are in before-tax real dollars. Different forms of income attract different tax rates, so the desired retirement income must be specified in before-tax dollars rather than after-tax dollars of spending. The average tax rate is almost always lower after retirement, especially after age 65. To set up the parameters for the model, you determine how much pre-tax income would be needed in retirement to provide the desired level of after-tax income that can be spent. You handle pensions by deducting their value from the required before-tax income in retirement.

Recall from the sections describing the models that we adopt a simple two-asset model for investment: equity and a risk-free asset. The equity return and standard deviation encompass all risky investment assets, including risky debt. The risk-free asset is a federal government treasury bill. We would include long-term government bonds in the risky asset class because of their interest rate risk.

Let us show how our model contrasts with the simple case study at the start of the paper. We set the desired probability that the amount accumulated for retirement will not be enough in retirement at no more than 5%. The probability of the entire plan succeeding will be worse than 5% because there is also the risk that the savings pre-retirement will not meet the required goal at retirement date. You could think of this as a recursive procedure or working backward. We first solve for the SPV of consumption in retirement that will give no more than a 5% probability of the person running out of money during retirement. Then, we give a savings plan and measure the probability of that plan meeting the required goal in the form of the SFV for the retirement date. Although it sounds like separate procedures, the model does all this in one pass.

The model does much more than the simple deterministic case and requires more inputs. For the pre-retirement savings parameters, the model adds a savings drift, volatility, and correlation between the savings rate and the risky asset's rate of return. The drift is somewhat analogous to the growth rate in the deterministic model, but it is not constant; it is a variable that, on average, is positive but varies since it has volatility. The correlation between the savings rate and the cumulative return is unique to Tahani & Robinson (2010) and this paper. If a family sees its portfolio is doing very well, it may be tempted to consume a bit more and save a bit less. So we expect the correlation to be either a small negative value or zero if the better returns don't influence savings behaviour.

The post-retirement consumption model has drift and volatility of consumption, correlation of consumption with the risky asset's rate of return, retirement age, and median age of death. The drift measures whether, on average, the annual real rate of consumption during retirement is likely to remain constant, rise, or fall. Other than Robinson & Tahani (2010) and this paper, other research has assumed a continuous consumption in retirement. The wealth effect may cause the correlation between consumption and the return on the risky asset to be slightly positive or zero if better returns do not induce increased consumption.

We start the demonstration with the stochastic equivalent of the simple case study. Recall that the model optimises the asset allocation between equity and a riskless asset to obtain the highest possible success probability. The portfolio return of 5.62% p.a. in the deterministic case study is specifically the result of our stochastic model's optimisation. We used it to make the deterministic and stochastic results as comparable as possible. Here are the inputs and the results of the stochastic model application to the case study:

- Equity returns 6% p.a., standard deviation 16%, riskless rate 2% p.a.
- Consumption in retirement parameters:
  - Consumption drift 0, Volatility 5%
  - Correlation (cumulative return, consumption) 10%
  - Initial consumption $50,000 p.a.
  - Retirement age 65; median age of death 85
  - Acceptable probability of ruin, i.e., running out of money, 5%
- Savings parameters:
  - Savings drift 2%, volatility 10%
  - Correlation (cumulative return, savings) -10%
  - Saving per year in first year $15,000
  - Initial wealth (that could be used for retirement) 0
  - Time horizon 35, i.e., start saving at age 30, retire at age 65

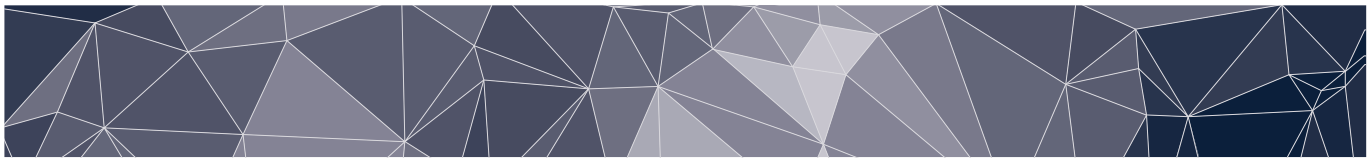

- Model results, rounded for easier reading:
  - Probability the plan will succeed: 48%
  - Expected wealth at 65: $2.1 million
  - Optimal asset allocation: 90% equity
  - Target wealth to allow $50,000 consumption with 5% risk: $1.8 million
  - The optimal consumption rate is 2.8% in retirement, i.e., $50,000 ÷ target wealth.

Target wealth is the amount the savings should accumulate to so that the retiree can consume $50,000 p.a. with only a 5% chance of running out of money.

Compare this with the deterministic result which gives expected wealth of $1.98 million and any reasonable fixed rate discount of future consumption even to age 100 is far less than $1.98 million. Apply Bengen's 4% rule and you get consumption in retirement of $79,000 p.a. which is far above the required $50,000 p.a. The deterministic procedure ignores the substantial variance risk during both savings and retirement phases. Bengen and others are using historical rates of return in their simulations for the retirement phase and the rates we used are comparable, but we allow for variance in return. The finding that an initial draw of 4% may not be sustainable in Finke et al. (2013) and Drew & Walk (2015) is consistent with our model results, although they look only at the retirement phase while we treat the entire life cycle.

Our stochastic model has many moving parts. Which parts matter the most? They fall into three categories: variables that have a major impact but are not ones that the planner or the client can control a lot; variables that are controllable; and variables that don't have a significant impact.

The risky rate of return and its standard deviation and the risk-free rate of return are determined exogenously. The overall risk and return profile can be managed a bit with proper asset allocation, but in the long run, they are going to converge to rates like the ones we used, which can be seen in any table of long-run returns for developed countries and world portfolios. If we run the model with even modest changes, we get significant changes in the plan's success rate. Still, we would get similar significance with a simple deterministic model. Likewise, the initial wealth (which is zero in our case study), the dollar amount saved in the first year, and the income desired in retirement have a large impact on the success probability. Several variables have little effect on the success probability when reasonable values other than the ones in the base case are entered: all the volatilities and the correlation of the risky asset's return with saving and consumption. The four factors that really do matter and are within the control of the consumer are acceptable probability of shortfall during retirement, drift factor of consumption in retirement, drift factor of savings, and retirement age. The median age of death requires a separate discussion.

Table 1 displays a sample of outcomes when we vary the four factors that are significant and controllable. Each entry shows three results separated by a "/": Probability the plan succeeds /

Table 1. Success Probability in % with Different Assumptions Canadian Males

Acceptable Failure	Base Case	(S)* Drift 1%	(C)* Drift -1%	(C) Drift -2%	Retire 68	Retire 68 (C) Drift -1%
5%	48/2.8/90	40/2.7/99	60/3.5/82	71/4.2/73	63/3.1/81	74/3.7/73
10%	65/3.7/84	56/3.6/92	75/4.5/74	83/5.2/66	77/4.0/73	85/4.7/65
20%	82/5.1/74	74/5.0/83	88/5.8/65	93/6.6/58	90/5.3/64	94/6.0/57

Three values: Probability the plan succeeds/Optimal consumption as a % of target wealth/% invested in risky asset.

Optimal consumption as a % of target wealth / % invested in the risky asset. The three rows are the acceptable probabilities of shortfall during retirement set at 5% (the base case), 10%, and 20%. The columns vary the two drift factors and retirement age by reasonable values. Everything else is kept at the base case values already set out, except median age, which increases to 86 in the case of retirement at age 68.

Start with the upper left cell in Table 1, the base case already described. The plan has a 48% probability of succeeding, which means the retired person or family gets to consume $50,000 in real dollars each year in retirement. The target retirement amount under the plan allows only 2.8% p.a. consumption of that amount per year starting in the first year of retirement. The optimal investment allocation to optimise the probability of success is 90% for risky assets.

The following two cells in the column show a significant effect when the acceptable probability of failure of the plan in retirement is raised to 10% and 20%. To repeat what this means – this probability is not the entire plan optimised, but rather the probability the person is willing to accept that they will arrive at retirement age with the accrued value, and it will not support $50,000 p.a. consumption throughout retirement. A failure doesn't mean they have no money; it means they must consume at least $1 less than $50,000 p.a. A change to 10% raises the probability of the plan succeeding to 65%, with consumption at 3.7% of the target and 84% in the risky asset. A rise to a 20% acceptable failure rate improves the probability of success to 82%.

The next column shows how important it is to keep increasing savings, even if the initial amount is quite large relative to the desired retirement income. The only change is to reduce the real drift factor to 1% from the 2% assumed in the base case, but this reduces the success probability by eight or nine percentage points.

The next two columns show the large effect a small negative drift (i.e., on average, a small decline) in retirement consumption has on the probability of success. For example, suppose you accept a 10% probability of a shortfall in retirement and expect consumption to decline by 1% in real dollars each year after you retire. In that case, the probability of success rises to 75%, and the risky asset allocation drops to 74%. This negative drift occurs after the transition to retirement, and the usual drop in consumption occurs when a person retires. Vettese and Morneau (2013, pg. 85 – 91) present a compelling argument that, on average, spending will decline during the retirement years as a person or couple becomes less active. They identify three phases of retirement activity and spending needs. Robinson and Tahani (2010) characterise three different behavioural patterns of spending during retirement. The 'gardener' has a negative consumption drift, the 'socialite has a constant real consumption with zero drift and the 'uninsured' has a positive spending drift because of medical costs that are not covered by universal health care in his country. Empirical evidence on consumption changes during the retirement years is difficult to interpret. Fisher *et al.* (2005, unpaginated) concludes: "While we find that consumption-expenditures decrease by about 2.5 percent when individuals retire, expenditures continue to decline at about a rate of 1 percent per year after that." They cite other papers with estimates of the decline ranging much more widely and discuss many issues that make all the estimates of declining expenses in retirement in their work and others challenging to rely on. We model the retirement plan of one person or family, and the appropriate drift factor in retirement will vary depending on the subject. Given the long-time frames, a family would be unlikely to have a good estimate and so with this evidence, a planner could reasonably use a drift factor for retirement consumption of -1%.

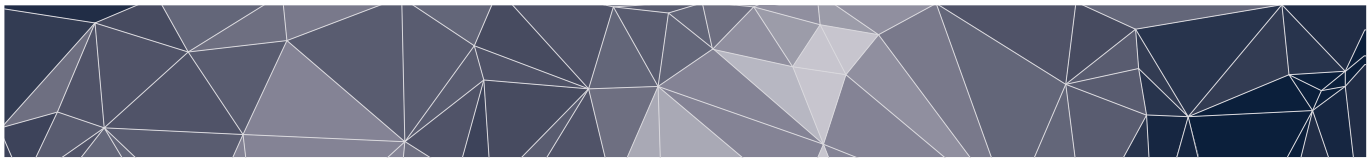

The next column shows the effect of a common plan for financing retirement. If you retire just a few years later, you increase the time to save and reduce the years of retirement to be financed. We chose 68 as the retirement age, which means 38 years of saving and a one-year increase in the median age of death to 86. The effect is considerable, though slightly less than the effect of a negative one per cent consumption drift. Combine working another three years with allowing a 20% probability of shortfall in retirement, and the success rate of the entire plan jumps to 90%. The last column shows a combination of retirement age 68 and a consumption drift in retirement of negative one per cent, yielding even higher probabilities of success.

The strategies of consumption now versus consumption in retirement and retirement later to allow more savings are well-established using deterministic planning. Our stochastic model allows a more informative estimate of the probability of success of such plans. Note also that even for the 94% success rate in the lowest right-hand corner for retirement at 68 and negative one per cent consumption drift, the allocation to the risky asset is 57% throughout the life cycle. Other allocations range up to 99%, with lower success rates. While the conclusions of any retirement model depend on the assumptions, the rates we have assumed in our numerical example are consistent with long-run history. Our results in Table 1 suggest that the classic advice to move much more of the portfolio into fixed income at retirement is likely outdated, because of higher life expectancy since that advice was first formulated as a rule of thumb.

The median age of death needs a separate discussion. Actuarial tables are too complex for a detailed discussion in this paper. Let us explain four crucial issues with the model and statistics on life expectancy. First, our model uses a commonly accepted function to translate an age of retirement plus a median date of death into mortality probabilities for the entire retirement period. The median date of death at any age is a conditional mortality calculation, which is different from the often-reported life expectancy at birth value for different countries. The life expectancy at birth statistic does not vary a lot among the developed world countries, from age 78 to 84 for both genders combined. For example, it is 78.5 in the US, 82.2 in Canada, 83 in Australia, and 84.2 in Japan (World Health Organization, 2023). We did not calculate the median age of death conditional on age 65 for all of them, but it is a few years older and has about the same difference across countries.

Second, our case study starts with a median age of death of 85, which is the figure for Canadian males aged 65. The comparable figure for Canadian females is 88 and 89 if she retires at age 68. Table 2 displays the success probabilities for the same combinations of assumptions as Table 1. For each combination the success rate for a Canadian woman is two to four percentage points lower; the most frequent value is four percentage points. This effect exacerbates the effect of lower average earnings of women during their working years, and the gender wage gap is universal (OECD, nd).

Table 2. Success Probability in % with Different Assumptions Canadian Females: Median Age of death 88 if retired at 65

Acceptable Failure	Base Case	(S) Drift 1%	(C) Drift -1%	(C) Drift -2%	Retire 68	Retire 68 (C) Drift -1%
5%	45	37	57	69	59	70
10%	61	52	72	81	73	82
20%	78	70	86	91	87	92

(*) Notation: Savings (S), Consumption (C).

Third, the life expectancy numbers we cite are for the entire country's population. Individual expectations are different and materially affected by factors like ethnicity, income, wealth, education, and social status. Individuals also know more about their own health and their ancestors' longevity. Accordingly, the median age of death appropriate for this model depends on the plan's subject(s).

Fourth, the model works on a person's mortality. When the family is a couple, the problem is not tractable analytically for a closed-form solution. Almost always, they will die at different dates, and after the first death, the survivor will have somewhat lower expenses. Joint mortality plus a non-linear decline in spending cannot be encompassed in our model. The best choice is to use the average of the two people's median ages of death.

When the chance of success is quite low, an interesting effect appears that we did not see in the tables. Take the base case but reduce the savings p.a. to $10,000 in the first year. The optimal equity allocation is 111%, which means the person gets the best shot at meeting the retirement goal by leveraging the risk-free rate. Since this only yields a 27% chance of success, we know that the downside is not only likely failure to have enough at retirement, but the person could fall disastrously short. Furthermore, the way the model works, this leverage is for life, not just pre-retirement. The proper advice the planner gives in this situation is not to lever the client but to tell the client that the desired retirement goal is virtually impossible to achieve at that level of saving.

5. Applying the Model in Practice

A competent derivatives practitioner can create a spreadsheet that incorporates this model to give a client a reasonable estimate of whether a proposed saving plan started today will allow the family to meet its retirement goal and what might need to be done to improve the risk. The usual possibilities can be canvassed: less consumption in retirement, more saving now, retiring later, and asset allocation.

In practice, the planner would enter values suitable to the individual client and the country in which they live. An important part of the discussion would be the determination of the level of a priori risk of shortfall in retirement since we have shown it has a large impact on the success probability. The life expectancy we used is Canadian, but that is not materially different from many other developed countries like Australia, New Zealand, and Japan. The difference between Canada and the US is enough that using its median life expectancy at age 65 would increase the success rates in Tables 1 and 2 by approximately two to four percentage points. The rates of return we used included Canada, US, and international equity, but a planner might choose different rates in different countries. We did not allow material fees to an investment manager, which would reduce success rates materially.

A planner doesn't run a model of a retirement plan once, and that is the end of the process. The world doesn't unfold as the plan proposed, and it must be revisited regularly. Our model gives a straightforward and easy way to observe progress and evidence to guide decisions on changes to the plan.

Lachance (2012) finds that a family should defer retirement savings to a later period of higher income to maximize utility. In practical terms, most families do exactly that because they have

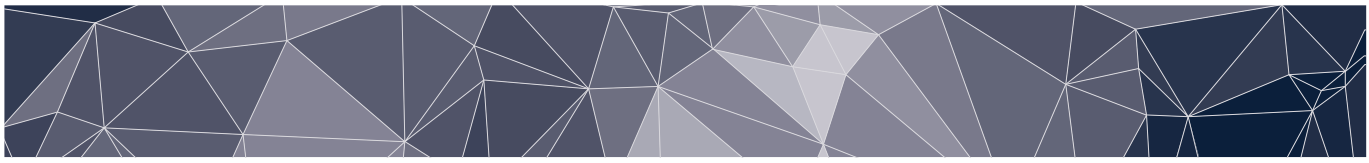

little choice with lower income, the cost of raising children, and buying a home. The planner would use our model at a later stage than our examples, starting at age 30, but the application would follow the same process.

Vettese & Morneau (2013) and Daley et al. (2018) point out that substantial wealth remains in owner-occupied housing that could also be liquidated to support consumption later in life. Home ownership is over 50% in most countries, and over 64% of families in Canada, Australia, the US, New Zealand, and the European Union own their own home. Our model and the other models cannot incorporate this directly. A practical way to give some idea of this value would be to annuitise the home's sale price at retirement and adjust the required income in retirement for the added cost of rent minus the costs of maintaining the home. This would significantly increase the success rate in most countries. A behavioural issue is that most people who own their own home do not want to leave it until much later in retirement.

We repeat our earlier cautions. This model uses pre-tax income because incorporating the complexities of different tax brackets and different tax rates on sources of income cannot be reduced to an analytical model with a closed-form solution. The results do not give a decision; they provide information for a client and a planner to use when making decisions on saving and investment. Too many factors are entering into the real-life situation that the model does not capture, nor do we know what future returns, changes in longevity, and behavioural changes in spending patterns will do for any given family and its retirement plan. For example, Lee and Hanna (2020, pg. 1) find: "Respondents with financial knowledge overconfidence (high subjective and low objective knowledge) were more likely to take early withdrawals than those with other combinations of objective and subjective knowledge."

Milevsky (2016) cautions against taking shortfall/success probabilities too literally. He suggests that a more straightforward way of explaining risk to most people would work better, given the difficulty most people already have with financial literacy, let alone with complex mathematical formulations. His suggestion, basically to get a reasonable estimate of how long your money needs to last for your lifetime and apply a reasonable real rate of return to the wealth at retirement, will yield a number that warns a family if they have a problem. Our more comprehensive model that covers the entire life cycle provides more helpful information, but it is hard to explain it all to most clients.

Collins et al. (2015) discuss the many problems of the various approaches to estimating the sustainability of a retirement income plan. They dismiss analytic approaches like ours because of the uncertainty of return distributions. However, any retirement plan suffers from these problems. The classic financial planning model with fixed rates of return, an assumed single date of death, and perhaps a few different scenarios has even more seriously limited assumptions. Our analytic model, which explicitly incorporates risks of returns, date of death, and spending and saving behaviour, provides more reliable information.

6. Conclusion and Future Research Directions

We present a stochastic model that estimates the probability of success of a financial plan starting with pre-retirement saving and ending at the end of the retirement period. We model the rate of return, date of death, saving pre-retirement, and consumption post-retirement as

stochastic variables. The model also yields the optimal investment allocation between a risk-free and a risky asset. We explain how the model works and apply it to a simple case study. Using reasonable values for the variables applied to a simple case study, we show that the traditional deterministic financial plan with fixed values will understate the risk of the plan. We find that the variables that the consumer can control and that affect the success probability the most are the drift rate of saving before retirement, the drift rate of consumption after retirement, the acceptable probability that the desired consumption cannot be maintained in retirement, and the date of retirement. We reiterate that our model applies to a particular family or person in a country, and the values the user inputs are specific to a single case each time.

An advance in our work could model consumption as mean-reverting to overcome one weakness in the previous work, both our own and simulation models. We know that consumption is mean-reverting. For example, buy a car this year but then drive it for a number of years, and consumption reverts to the longer-run mean, whether that is slowly declining, rising, or staying constant in real dollars. If we don't make consumption mean-reverting, a certain percentage of the time, it will keep going up or going down indefinitely in the model, but in practice, that doesn't happen. We would also like to incorporate a mean-reverting rate of return since we know that it is closer to observed return behaviour over time. We have not determined if we can incorporate mean reversion into a closed form solution.

Estrada (2023) points the way to improving the information our stochastic model could provide. He produces percentiles of retirement consumption coverage ratios for different asset allocations between equity and bonds. We cover the entire life cycle, and our primary output is a probability of success, not a metric like coverage ratio. Perhaps we can extract from the model the deciles of consumption in retirement and the probability of reaching each decile conditional on an optimal pre-retirement saving plan, thus showing the user more information about what could happen outside the optimal case. This would require an entirely new paper, not just an added tweak to the current paper.

A final possibility that would also require another paper is flipping the model to output the required annual savings. Specify, say, a 20% probability of failing to reach the desired consumption in retirement and a 70% probability that the entire plan succeeds; then, how much saving is required each year?

References

Bajtelsmit, V & Wang, T 2018, 'Household financial planning strategies for managing longevity risk,' *Financial Planning Review,* vol. 1, nos. 1-2, e1007. https://doi.org/10.1002/cfp2.1007

Bengen, W 1994, 'Determining Withdrawal Rates Using Historical Data,' *Journal of Financial Planning,* pp. 14-24.

Collins, P, Lam, H & Stampfli, J 2015, 'How risky is your retirement income risk model?' *Financial Services Review,* vol. 24, no. 3, pp. 193-216.

Daley, J, Coates, B, Wiltshire, T, Emslie, O, Nolan, J & Chen, T 2018, *Money in retirement: More than enough.* Grattan Institute. 142 pages. https://grattan.edu.au/wp-content/uploads/2018/11/912-Money-in-retirement.pdf

Drew, M & Walk, A 2015, 'Just how safe are 'safe withdrawal rates' in retirement?' *Financial Planning Research Journal,* vol. 1, pp. 22-33.

Estrada, J 2023, 'Retirement Planning—Is One Number Enough?' *Journal of Financial Planning*, vol. 36, no. 7, pp. 64-81.

Estrada, J & Kritzman, M 2019, 'Toward Determining the Optimal Investment Strategy for Retirement,' *Journal of Retirement*, Summer, pp. 35-42.

Finke, M, Pfau, W & Blanchett, D 2013, 'The 4 Percent Rule Is Not Safe in a Low-Yield World,' *Journal of Financial Planning*, vol. 26, no. 6, pp. 46-55.

Fisher, J, Johnson, D S, Marchand, J, Smeeding, T & Torrey, B 2005, 'The Retirement Consumption Conundrum: Evidence from a Consumption Survey' Working Paper 2005-14.' Chestnut Hill, MA: *Center for Retirement Research at Boston College*. https://crr.bc.edu/the-retirement-consumption-conundrum-evidence-from-a-consumption-survey/

Ho, K, Milevsky, M A & Robinson, C 1996, 'Risk-adjusted retirement.' *Canadian Investment Review*, Spring, pp. 19-27.

Ho, K, Robinson, C, Letkiewicz, J & Zaremba, V 2022, *Personal Financial Planning*, Captus Press.

Huang, H, Milevsky, M A & Salisbury, T 2017, 'Retirement spending and biological age,' *Journal of Economic Dynamics & Control,* vol. 84, pp. 58-76.

Lachance, M E 2012, 'Optimal onset and exhaustion of retirement savings in a life-cycle model,' *Journal of Pension Economics and Finance*, vol. 11, no. 2, pp. 21-52.

Lee S T, & Hanna S 2020, 'Financial knowledge overconfidence and early withdrawals from retirement accounts,' *Financial Planning Review*, Vol. 3, no. 2, e1091. https://doi.org/10. 1002/cfp2.1091

MacDonald, B J, Morrison, R J & Avery, M 2018, *Exploring The 'Risk' of Ignoring Risks in Retirement Financial Planning*, Working paper. https://fpcanadaresearchfoundation.ca/media/g2nfiwsj/understanding-the-importance-of-risk-for-financial-planning.pdf

Milevsky, M A 1997, 'The present value of a stochastic perpetuity and the Gamma distribution,' *Insurance: Mathematics and Economics,* vol. 20, pp. 243-250.

Milevsky, M A 2016, 'It's Time to Retire Ruin (Probabilities),' *Financial Analysts Journal,* March/April, pp. 8-12.

Milevsky, M A & Robinson, C 2000, 'Self-annuitization and ruin in retirement,' *North American Actuarial Journal,* vol. 4, pp. 1-17.

Milevsky, M A & Robinson, C 2005, 'A Sustainable Spending Rate Without Simulation,' *Financial Analysts Journal,* vol. 61, no. 6, pp. 89-100.

OECD nd, *Gender Wage Gap*, https://data.oecd.org/earnwage/gender-wage-gap.htm

Robinson, C & Tahani, N 2010, 'Sustainable Retirement Income for the Gardener, the Socialite and the Uninsured,' *Financial Services Review*, vol. 19, no.3, pp. 187-202.

Suarez, E D 2020, 'The perfect withdrawal amount over the historical record,' *Financial Services Review*, vol. 28, no. 2, pp. 96-132.

Suarez, E D, Suarez, A & Walz, D 2015, 'The perfect withdrawal amount: A methodology for creating retirement account distribution strategies,' *Financial Services Review*, vol. 24, no. 4, pp. 331-57.

Tahani, N & Robinson, C 2010, 'Freedom 55 or Drudgery until 70?' *Financial Services Review*, vol. 19, no. 4, pp. 275-84.

Vettese, F & Morneau, B 2013, *The Real Retirement*, John Wiley.

Willows, G 2020, 'Develop a retirement plan and stick to it: it will improve both your attitude and behaviour with money,' *Financial Services Review*, vol. 20, no. 3, pp. 243-71.

World Health Organization 2023, *Country, area, WHO region and global health statistics, spreadsheet*, https://www.who.int/data/gho/publications/world-health-statistics

Yaari, M 1965, 'Uncertain lifetime, life insurance, and the theory of the consumer,' *The Review of Economic Studies,* vol. 32, no. 2, pp, 137-150.

Appendix A

We define the *Stochastic Future Value* (SFV) and compute its first two moments. We then show how we can approximate its probability distribution. Given the cumulative return process (R_t) in Eq. (4), the future value at time T of the initial wealth w is simply wR_T. On the other hand, an investment of S_u at time u is equivalent to a present value of $\frac{S_u}{R_u}$ at time 0; or to a future value of $R_T \frac{S_u}{R_u}$ at time T. Since the savings are invested continuously, their total future value is obtained by integrating over the time interval $[0, T]$. The SFV is then obtained as in Eq. (1) by:

$$SFV = wR_T + \int_0^T S_u \frac{R_T}{R_u} du \tag{A.1}$$

As seen in Eq. (4), both processes (R_t) and (S_t) are Geometric Brownian motions. Therefore, the SFV can be seen as a (continuous) sum of lognormal variables. Thanks to the literature on Asian options (i.e. arithmetic average options), many moment-matching-based approximations with known distributions are available.

Combining Eq. (1) and Eq. (4), the SFV can be written as:

$$SFV = we^{\bar{\theta}T + \tilde{\sigma}B_T} + se^{\bar{\theta}T} \int_0^T e^{(\bar{\eta}-\bar{\theta})u} e^{\tilde{\sigma}(B_T - B_u)} e^{\beta Z_u} du \tag{A.2}$$

Using standard stochastic calculus, it can be shown that the first two moments of the SFV are given by:

$$\begin{cases} E_1 = e^{\tilde{\mu}T} \times \left(w + \frac{s}{A_1}\left(e^{A_1 T} - 1\right)\right) \\ E_2 = e^{A_2 T} \times \left(w^2 + \frac{2ws}{A_3}\left(e^{A_3 T} - 1\right) + \frac{2s^2}{A_4 A_5}\left(e^{A_4 T} - 1\right) - \frac{2s^2}{A_3 A_5}\left(e^{A_3 T} - 1\right)\right) \end{cases} \tag{A.3}$$

where $(A_i)_{i=1,\ldots,5}$ are functions of the allocation parameter α given by:

$$\begin{cases} A_1 = \eta - (1-\alpha)r - \alpha\mu \\ A_2 = 2(1-\alpha)r + 2\alpha\mu + \alpha^2\sigma^2 \\ A_3 = \eta + \rho\alpha\beta\sigma - \alpha^2\sigma^2 - (1-\alpha)r - \alpha\mu \\ A_4 = 2\eta + \beta^2 - A_2 \\ A_5 = \eta + \beta^2 - \rho\alpha\beta\sigma - (1-\alpha)r - \alpha\mu \end{cases} \tag{A.4}$$

Matching the first two moments E_1 and E_2 with those of a lognormal distribution with parameters a and b lead to Eq. (6).

Appendix B

We show that the probability of ruin can be defined by the probability that the SPV is below the wealth at retirement W. Following Equations (2) and (8), we define the investment value (R_t) and the consumption (C_t) as two correlated Geometric Brownian motions given by:

$$\begin{cases} dR_t = \tilde{\mu} R_t dt + \tilde{\sigma} R_t dB_t \quad ; \quad R_0 = 1 \\ dC_t = -\kappa C_t dt + \psi C_t dZ_t \quad ; \quad C_0 = 1 \\ d < B, Z >_t = \rho' dt \end{cases} \tag{B.1}$$

where the correlation between the two Brownian motions is ρ'. Assuming that we can only invest in R and consume C per time unit, the wealth W_t at time t can be given as:

$$dW_t = \left(\mu W_t - C_t\right) dt + \sigma W_t dB_t \quad ; \quad W_0 = G \tag{B.2}$$

Using Itô's lemma, we can easily show that:

$$W_t = R_t \times \left[G - \int_0^t C_s R_s^{-1} ds \right] \tag{B.3}$$

solves the SDE (B.2). It is obvious that:

$$W_T \leq 0 \Leftrightarrow G \leq \int_0^T C_s R_s^{-1} ds \tag{B.4}$$

for any future time T. This shows that the probability of ruin prior to some time T is equal to:

$$\begin{aligned} P\left(\inf\left\{W_s; 0 \leq s \leq T\right\} \leq 0\right) &= P\left(W_T \leq 0\right) \\ &= P\left(G \leq \int_0^T C_s R_s^{-1} ds \right) \end{aligned} \tag{B.5}$$

where the first equality comes from the fact that the integral $\int_0^T C_s R_s^{-1} ds$ is non-decreasing with respect to T.