

Recreating Sustainable Retirement

Resilience, Solvency, and Tail Risk

EDITED BY

Olivia S. Mitchell,
Raimond Maurer, and
P. Brett Hammond

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Pension Research Council, The Wharton School, University of Pennsylvania 2014

The moral rights of the authors have been asserted

First Edition published in 2014

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2014940448

ISBN 978-0-19-871924-3

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>List of Figures</i>	ix
<i>List of Tables</i>	xiii
<i>Notes on Contributors</i>	xv
1. Recreating Retirement Sustainability <i>Olivia S. Mitchell and Raimond Maurer</i>	1
Part I. Capital Market and Model Risk	
2. Managing Capital Market Risk for Retirement <i>Enrico Biffis and Robert Kosowski</i>	9
3. Implications for Long-term Investors of the Shifting Distribution of Capital Market Returns <i>James Moore and Niels Pedersen</i>	30
4. Stress Testing Monte Carlo Assumptions <i>Marlena I. Lee</i>	60
Part II. Longevity Risk	
5. Modeling and Management of Longevity Risk <i>Andrew Cairns</i>	71
6. Longevity Risk Management, Corporate Finance, and Sustainable Pensions <i>Guy Coughlan</i>	89
7. Model Risk, Mortality Heterogeneity, and Implications for Solvency and Tail Risk <i>Michael Sherris and Qiming Zhou</i>	113
8. The Securitization of Longevity Risk and Its Implications for Retirement Security <i>Richard MacMinn, Patrick Brockett, Jennifer Wang, Yijia Lin, and Ruilin Tian</i>	134

viii Contents

Part III. Regulatory and Political Risk

9. Evolving Roles for Pension Regulations: Toward Better Risk Control? 163
E. Philip Davis
10. Developments in European Pension Regulation: Risks and Challenges 186
Stefan Lundbergh, Ruben Laros, and Laura Rebel
11. Extreme Risks and the Retirement Anomaly 215
Tim Hodgson

Part IV. Implications for Plan Sponsors

12. Risk Budgeting and Longevity Insurance: Strategies for Sustainable
Defined Benefit Pension Funds 247
Amy Kessler
13. The Funding Debate: Optimizing Pension Risk within a Corporate
Risk Budget 273
Geoff Bauer, Gordon Fletcher, Julien Halfon, and Stacy Scapino
- The Pension Research Council* 293
- Index* 297

Chapter 4

Stress Testing Monte Carlo Assumptions

Marlena I. Lee

Monte Carlo simulations are a useful financial planning tool serving several purposes. They are often used to forecast wealth outcomes into the future for the purposes of financial planning.¹ One can input assumptions about returns, saving, and spending needs, and the simulation reports how likely these goals are to be achieved, given the assumptions of the model. This framework is immensely valuable for helping investors understand the key factors that can influence their long-term investment goals. But Monte Carlo simulations also require some key assumptions to simplify a very complex problem.

This chapter examines the importance of three assumptions central to most Monte Carlo simulations: that stock returns are normally distributed, expected returns are constant over time, and return parameters are known to the user. None of these assumptions is completely consistent with reality. Returns are known to have fat tails, meaning extreme events have occurred more often than expected under a normal distribution. Expected returns also likely vary through time and, most importantly, they are unobservable, let alone precisely measured. Although underlying Monte Carlo assumptions are not an exact description of the world, the usefulness of the simulations depends on whether they can still provide useful insights to guide financial planning. Given that we know the assumptions are only approximations, this chapter assesses how users might interpret Monte Carlo results that attempt to forecast wealth outcomes into the future.

The Behavior of Black Swans

Stock returns tend to have more extreme observations than would be expected under a normal distribution. For example, global market declines in excess of 40 percent should occur once every 540 years according to a normal distribution.² Yet since 1900, global markets have seen two such declines, once in 1931 during the Great Depression, and again in 2008 during the global financial crisis. Researchers have studied these tail returns since the early 1960s, although popular interest in these ‘black swans’ has risen with the recent bestselling book of the same name.³

Figure 4.1 shows a histogram of global equity returns relative to the normal distribution. This is generated using annual returns on global equities from 1900

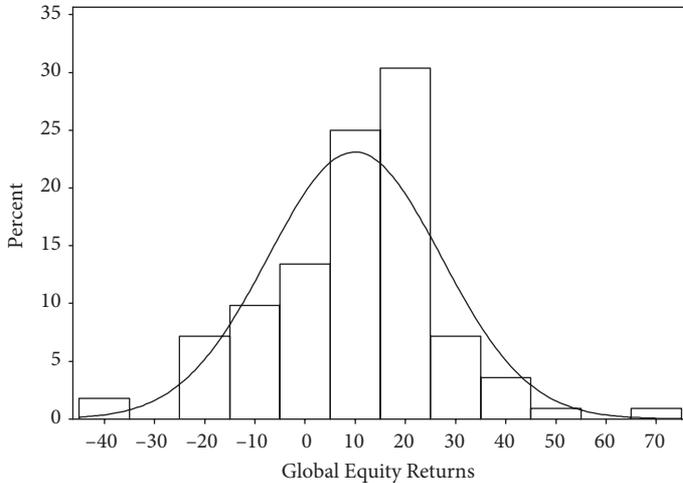


Figure 4.1. Histogram of annual equity returns: 1900–2011.

Source: Dimson et al. (2011).

to 2011 from the Dimson, Marsh, and Stauton Global Returns Database (Dimson et al. 2011). Much attention is often devoted to the negative tail events, although Figure 4.1 shows that extreme positive events also occur, such as the 71 percent return in 1933. Returns also tend to cluster around the mean much more than expected under a normal distribution. These patterns tend to offset, so that the assumption of normality does not greatly impact the results of Monte Carlo simulations.

To illustrate this point, 100,000 outcomes of a 100 percent global equity portfolio are simulated for 30 years. In the baseline Monte Carlo simulation, returns are assumed to be normally distributed with mean and standard deviation equal to that observed in historical data. A second simulation ‘bootstraps’ returns from actual historical data. This bootstrap simulation preserves many of the characteristics of the actual distribution including the mean, standard deviation, fat tails, and skewness of returns. This means that market drops of 40 percent will occur 1.79 percent of the time, instead of 0.19 percent of the time. We repeat each simulation 100,000 times.

Results are shown in Table 4.1. Historical returns appear in Row A, and the simulated returns in Rows B1–B2. As expected, both simulations produce averages and standard deviations t which are very similar to those observed in the historical data. Annualized average returns are within a few basis points, an acceptable range for sampling error. The bootstrap methodology better captures specific percentiles of the historical distribution. Tail percentiles and the median match up

62 Recreating Sustainable Retirement

TABLE 4.1 Summary statistics of global equity returns, 1900–2011

	Avg.	Std. dev.	Percentiles					N
			5th	10th	50th	90th	95th	
A. Historical returns	9.91	17.26	-18.09	-13.11	11.45	28.23	36.38	112
B. Simulated returns								
B1. Baseline Monte Carlo	9.93	17.27	-18.50	-12.21	9.92	32.05	38.31	3,000,000
B2. Bootstrap	9.97	17.20	-18.09	-13.11	11.73	28.23	36.38	3,000,000
B3. Bootstrap 10yr	9.92	17.19	-18.09	-13.11	11.73	28.23	36.38	3,000,000
B4. Random mean	9.90	17.34	-18.64	-12.34	9.91	32.12	38.41	3,000,000
C. Difference in annualized return from historical								
C1. Baseline Monte Carlo	0.02		-0.41	0.90	-1.53	3.82	1.93	
C2. Bootstrap	0.07		0.00	0.00	0.28	0.00	0.00	
C3. Bootstrap 10yr	0.01		0.00	0.00	0.28	0.00	0.00	
C4. Random mean	-0.01		-0.55	0.77	-1.54	3.89	2.03	

Source: Author's calculations derived from Dimson et al. (2011).

almost perfectly in the bootstrap, while small discrepancies appear when returns are simulated from a normal distribution.

Table 4.2 shows growth of wealth statistics after 30 years.⁴ The baseline Monte Carlo simulation, in Row A1, produces very similar results to the bootstrap simulation in Row A2, although with slightly narrower tails. In other words, fat tails in the return distribution also result in fat tails in the wealth distribution over time. Annualized returns at the fifth percentile differ by 22 basis points per year. Accordingly, the assumption that returns are normally distributed does not greatly impact the results of Monte Carlo simulations.⁵

Time-varying Expected Returns

Both theoretical and empirical research indicates that the expected equity risk premium varies over time. Investors demand a higher expected return to hold risky assets such as stocks, and this expected return should be higher during times of greater uncertainty, such as at the onset of a recession. Some argue that this results

TABLE 4.2. Simulated growth of \$1

	Average	Std dev	Percentiles				
			5th	10th	50th	90th	95th
A. Growth of \$1 over 30 years							
A1. Baseline Monte Carlo	15.55	15.74	2.51	3.47	10.89	32.63	44.02
A2. Bootstrap	15.79	15.97	2.35	3.36	11.03	33.29	45.10
A3. Bootstrap 10yr	14.59	14.13	3.15	3.97	10.21	29.85	40.88
A4. Random mean	15.44	15.78	2.44	3.41	10.74	32.50	43.73
B. Difference in annualized returns from baseline Monte Carlo (basis points)							
B1. Baseline Monte Carlo	–	–	–	–	–	–	–
B2. Bootstrap	6		–22	–11	4	8	9
B3. Bootstrap 10yr	–23		79	48	–23	–33	–28
B4. Random mean	–3		–9	–6	–5	–2	–3

Source: Author’s calculations derived from Dimson et al. (2011).

in mean reversion in returns: following periods of poor market returns, expected returns should be higher than average.⁶

All time series patterns in stock returns such as autocorrelation or variation around business cycles are assumed away in a typical Monte Carlo simulation. Each year, expected returns are assumed to be the same, regardless of economic conditions or recent returns. It turns out that this is not a bad assumption. Using the framework from the previous section, next we show that incorporating time series patterns in returns has only a mild impact on simulation results.

We simulate a 100 percent global equity portfolio using a bootstrap simulation very similar to that in the previous section, but with one important difference. Instead of randomly selecting annual returns, we now randomly select ten-year returns by picking a random year t between 1900 and 2011. The return of a simulated portfolio for its first ten years equals the return in years t to $t + 9$.⁷ We repeat the process two more times, until we have returns for a 30-year investment horizon. Row B3 in Table 4.1 shows that the simulated returns are very similar to the actual distribution of returns. The benefit of this method is that it captures any variation in expected returns that might occur over the course of ten-year return cycles.

This simulation generates portfolio outcomes with lower averages and less wealth dispersion than the baseline Monte Carlo simulation, as shown in Row A3 of Table 4.2. Tighter wealth outcomes are the result of very slight mean reversion in the global equity data. Average returns are slightly higher following periods of

64 Recreating Sustainable Retirement

poor returns. This makes long-run returns, such as the ten-year returns in the simulations, less risky than ten independently drawn annual observations.⁸ The effect is that very bad and very good outcomes are less likely. In the lower tail at the fifth percentile, this reduction in long-term risk works out to a positive difference in annualized returns of 79 basis points. This benefit comes at the cost of lower annualized returns in the upper part of the wealth distribution.

Although mean reversion is good news for long-run investors, investors should be cautious about whether mean reversion will also occur in the future. Most of the mean reversion in the global portfolio comes from the U.S., and results using World ex-U.S. do not display evidence of mean reversion. Mean reversion is difficult to assess since even with more than a century of returns, the sample contains only 11 completely independent ten-year return observations. Given uncertainty about whether these patterns will hold in the future, one sensible approach would be to assume that expected returns are constant through time. If mean reversion did appear in the future, this assumption will prove to be slightly conservative.

Uncertain Inputs

All of the simulations thus far make a critical assumption, namely that the true return distribution is known. Of course it is not. Nominal global equity returns have averaged about 10 percent over the period from 1900 to 2011, but there is substantial uncertainty about future expected returns.⁹ Even if the distribution has not changed over time, there is a good chance that the expected nominal equity return ranges anywhere from 7 to 13 percent.¹⁰

To examine the impact of uncertainty around expected returns, next we run a two-step simulation. We first randomly draw an expected return, $\tilde{\mu}$, from a normal distribution with mean equal to the historical average and standard deviation equal to the standard error.¹¹ Next, we randomly draw a return from a normal distribution with random mean $\tilde{\mu}$ and standard deviation equal to the sample standard deviation. This simulation accounts for uncertainty around the expected return, but still assumes the standard deviation is known without error.

Return summary statistics in Row B4 of Table 4.1 show that this simulation produces a very similar return profile to the other simulations. The extreme tails of the distribution are slightly wider than the baseline Monte Carlo simulation, and because it assumes returns are normally distributed, it does not perfectly match all points of the historical distribution.

Wealth outcomes in Row A4 of Table 4.2 show essentially identical results to the baseline Monte Carlo simulation. As long as the historical average is an unbiased estimate of the true mean, which is true if the distribution of returns does not change over time, then uncertainty about the mean has no impact on the simulation. But it is unclear whether historical averages can be considered an unbiased estimate of future expected returns.

TABLE 4.3. Monte Carlo simulations with varying expected returns

		Avg.	Std dev	Percentiles				
				5th	10th	50th	90th	95th
C =	-2%	9.13	9.48	1.41	1.96	6.31	19.31	26.18
	-1%	11.93	12.23	1.88	2.61	8.30	25.14	33.99
	0%	15.55	15.74	2.51	3.47	10.89	32.63	44.02
	1%	20.22	20.21	3.33	4.59	14.26	42.27	56.89
	2%	26.24	25.90	4.41	6.06	18.61	54.64	73.33

Note: Growth of \$1 over 30 years; expected return = historical average + C.

Source: Author’s calculations derived from Dimson et al. (2011).

The average historical excess return of U.S. stocks over one-month Treasury bills has been about 8 percent from 1926 to 2011. Some Monte Carlo users may also use this figure as their assumption for the expected equity premium going forward. But there are reasons to believe this estimate might be high. Extending the sample back to 1900 yields a U.S. equity premium of about 7 percent, and expanding globally lowers the equity premium to about 6 percent. Using long-term dividend and earnings growth, Fama and French (2002) estimate an equity premium in the range of 2.6 to 4.3 percent. If future expected returns are lower than historical averages, the impact on Monte Carlo simulations can be large.

Table 4.3 shows wealth outcomes for standard Monte Carlo simulations but with different levels of expected return. The simulation with mean equal to the historical average embeds an equity premium of about 6 percent. If the equity premium is actually 8 percent, the simulations in the first row (C = -2 percent) would apply. Similarly, the last row corresponds to simulations with expected return of about 8 percent over T-bills. The results are dramatically different, with each percentage difference in expected return cumulating to a 30 percent difference in wealth over 30 years. These results dwarf the minor deviations that result from fat tails and mean reversion, and they highlight a critical flaw in Monte Carlo simulations when users use an upward-biased expected return assumption.

Conclusion

Monte Carlo simulations incorporate many assumptions that simplify reality. These assumptions are not perfect descriptions of the world, but they appear to be decent approximations for some purposes. Moreover, simulation methods that better reflect historical returns do not dramatically impact results in our setting. Bootstrapping returns to account for extreme tail returns has little impact on the simulations relative to a simple assumption that returns are normal. And although

66 Recreating Sustainable Retirement

expected returns on equities do vary through time, it seems reasonable to simply assume that expected returns are constant through time.

One important assumption that does have a critical impact when using Monte Carlo simulations to project absolute future wealth is the long-run expected rate of return on equities. Changing expected return assumptions dwarfs differences that arise from all other assumptions examined in this study. When using Monte Carlo to project future wealth, no tool, no matter how many bells and whistles, can escape this fundamental problem. Expected future returns are unobservable and are incredibly difficult to estimate precisely.

In our view, Monte Carlo simulations are a very useful financial planning tool. But understanding the tool's limitations will help prevent its misuse. Monte Carlo output cannot be interpreted as a guarantee, since the model does not account for a myriad of factors that can impact investment outcomes. Instead, Monte Carlo simulations should be viewed as a directional guide to let investors know if they are roughly on track. Combined with frequent evaluation and investor discipline, Monte Carlo simulations are a useful component of a sound financial plan to help increase the probability of investor success.

Disclaimer

The projections or other information generated by Monte Carlo analysis tools regarding the likelihood of various investment outcomes are hypothetical in nature, do not reflect actual investment results, and are not guarantees of future results. Results may vary with each use and over time. These hypothetical returns are used for discussion purposes only and are not intended to represent, and should not be construed to represent, predictions of future rates of return. Actual returns may vary significantly.

Notes

1. This chapter studies the use of Monte Carlo simulations for predicting future wealth outcomes. Other uses, such as to assess liability hedging, are not analyzed in this chapter.
2. Computed using a mean of 9.9 percent and standard deviation of 17.26 percent, the sample estimates from annual global returns from 1900 to 2011.
3. See Mandelbrot (1963), Fama (1963), and Taleb (2010).
4. Investment horizons of ten and 20 years yield similar conclusions.
5. A greater impact would be observed as one moves more into the tails. I have only examined the fifth percentile here, although the differences in results may be more pronounced in the first percentile.
6. Jorion (2003) shows that the empirical evidence on mean reversion in historical stock returns is weak, particularly in global returns.
7. For $t \geq 2003$, I use returns from t to 2011, then from 1900 on until I have a ten-year period. This is to ensure that returns in the first and last ten years are not under-sampled.

8. This is not to say that long-term returns are less risky than short-term returns. The distribution of wealth outcomes grows with the investment horizon.
9. It is more common to assume an expected equity premium over a risk-free rate. Since my goal is to illustrate the impact of uncertainty around expected returns, I examine total equity returns for the sake of simplicity.
10. If returns are normally distributed, a 95 percent confidence interval for the mean would range from 6.7 percent to 13.1 percent.
11. The standard error equals the standard deviation divided by \sqrt{N} .

References

- Dimson, E., P. Marsh, and M. Staunton (2011). *The Dimson-Marsh-Staunton Global Investment Returns Database*. New York: Morningstar, Inc.
- Fama, E. F. (1963). 'Mandelbrot and the Stable Paretian Hypothesis,' *Journal of Business*, 36: 420–249.
- Fama, E. F., and K. R. French (2002). 'The Equity Premium,' *Journal of Finance*, 57: 637–659.
- Jorion, P. (2003). 'The Long-Term Risks of Global Stock Markets,' *Financial Management*, 32: 5–26.
- Mandelbrot, B. (1963). 'The Variation of Certain Speculative Prices,' *Journal of Business*, 36(4): 394–419.
- Taleb, N. N. (2010). *The Black Swan: The Impact of the Highly Improbable*. New York, NY: Random House.