



Kai Philipp Winselmann

Essays on the Kelly Criterion and Growth Optimal Strategies

Dissertation presented to the head of the Doctoral Committee

in partial fulfillment of the requirements for the degree of

Doctor rerum politicarum (Dr. rer. pol.)

in the subject of economic science

WHU – Otto Beisheim School of Management

Vallendar, Germany

January 2018

Referee: Prof. Dr. Markus Rudolf

Co-Referee: Prof. Dr. Stefan Spinler

Acknowledgements

The present doctoral thesis was written during my time as research assistant and doctoral candidate at the Endowed Chair of Finance at WHU – Otto Beisheim School of Management. I will remember this time as an exceptionally insightful and enjoyable period, both privately and professionally.

First and foremost, I would like to thank my supervisor Prof. Dr. Markus Rudolf, who greatly inspired my interest for the topic of this thesis and at the same time granted me the academic freedom to pursue research questions and methodologies that I am passionate about. Simultaneously, he challenged my thinking and research at all stages of the dissertation project and thereby fostered my own deeper understanding of the subject and the research questions to answer. Likewise, I owe a debt of gratitude to my co-supervisor Prof. Dr. Stefan Spinler for his academic and personal guidance throughout my time at WHU.

For their advice, time, and patience, I am very much beholden to Asst.-Prof. Dr. Julia Kapraun and my colleagues at the Endowed Chair of Finance. Not only did they support my academic pursuits, but they also made for many joyful days and evenings filled with discussions, laughter and friendship. Moreover, I wish to thank Bain & Company for granting me an educational leave that allowed me to take this exciting excursion into academic research while still remaining part of the firm at all times.

Last but not least I deeply thank my family for their support and encouragement of all my academic, professional and private endeavors. The learning, the memories, and the benevolence of those who accompanied me throughout the years will be a source of enduring happiness.

Kai Philipp Winselmann
Munich, January 2018

Contents

List of Figures	IV
List of Tables	V
List of Symbols and Abbreviations	VI
1 Introduction	1
2 Growth Optimal Strategies and the Kelly Criterion: A Review of Theory and Applications	3
2.1 Introduction	3
2.2 The Kelly criterion	4
2.2.1 Computation of the growth optimal wager	6
2.2.2 Expected utility theory foundations	8
2.2.3 Fractional Kelly strategies	11
2.2.4 Critique of the Kelly criterion and logarithmic utility	12
2.3 Applications	14
2.3.1 Portfolio choice and asset pricing	14
2.3.2 Risk and the Kelly criterion	18
2.3.3 Kelly strategies in gambling and investing	20
2.4 Conclusion	24
3 Using the Kelly Criterion to Protect Wealth Growth in Myopic Multi-Period Mean-Variance Portfolio Selection	25
3.1 Introduction	25
3.2 Asset allocation decisions	28
3.2.1 Growth optimal investor	30
3.2.2 Mean-variance investor	33
3.3 Theoretical considerations	35
3.3.1 Overbetting	35
3.3.2 The weight in the risky asset	37
3.3.2.1 Growth optimal investor	37
3.3.2.2 Mean-variance investor	38
3.3.2.3 Comparison of risk taking preferences	40

3.4	Application example	43
3.4.1	Constant investment opportunities	43
3.4.2	Stochastic investment opportunities	48
3.5	Conclusion	52
4	A Performance Comparison of Growth Optimal Strategies and Myopic Mean-Variance Optimization for 80 Years of Stock Market Investing	54
4.1	Introduction	54
4.2	The model	58
4.2.1	Mean-variance investor	59
4.2.2	Growth optimal investor	60
4.3	Data and methodology	63
4.3.1	Data	63
4.3.2	Parameter estimates	64
4.3.3	Quality of estimates	65
4.4	Results	66
4.4.1	Investor wealth paths	67
4.4.2	Robustness to variations in ϕ and γ	71
4.4.3	Robustness to variations in period length	74
4.4.4	Robustness to variations in estimation period length	75
4.4.5	Fractional Kelly strategies	77
4.5	Conclusion	78
4.A	Appendix	80
5	Conclusion	83
	Bibliography	VIII

List of Figures

2.1	Binomial tree depicting the wealth development for two sequential Bernoulli trials	6
3.1	Investor weight in the risky asset for various combinations of risky asset mean excess return and standard deviation	39
3.2	Surface plot of investor weight in the risky asset for various combinations of risky asset mean excess return and standard deviation	40
3.3	Overbetting of mean-variance investors with four different values of risk aversion γ	42
3.4	Average final wealth and Sharpe ratios of growth optimal and mean-variance investors with constant investment opportunities	45
3.5	Average final wealth and Sharpe ratios of growth optimal and mean-variance investors with stochastic investment opportunities	50
4.1	Histogram of deltas between actual return and estimated mean return with normal distribution graph	66
4.2	Wealth paths of selected investors	68
4.3	Final wealth and Sharpe ratios for varying values of ϕ and γ	72
4.A.1	Final wealth and Sharpe ratios for bi- and trimonthly re-balancing	80

List of Tables

3.1	Weight of the risky asset in mean-variance investor portfolios	44
3.2	Simulated final wealth and bankruptcies of mean-variance investors and growth optimal investors under constant investment opportunities	46
3.3	Simulated final wealth and bankruptcies of mean-variance investors and growth optimal investors under stochastic investment opportunities	49
4.1	Descriptive statistics of data set	65
4.2	Significance of differences in returns over time	69
4.3	Performance measures for selected investors	70
4.4	Final wealth levels and Sharpe ratios for various values of ϕ and γ . .	73
4.5	Final wealth levels and Sharpe ratios for bimonthly re-balancing . . .	75
4.6	Final wealth levels and Sharpe ratios for trimonthly re-balancing . .	75
4.7	Final wealth levels and Sharpe ratios for monthly re-balancing and various estimation period lengths	76
4.8	Final wealth levels and Sharpe ratios with fractional Kelly strategies	78
4.A.1	Arithmetic performance measures for high growth investors in five sub-periods	81
4.A.2	Arithmetic performance measures for equal growth investors in five sub-periods	81
4.A.3	Geometric performance measures for high growth investors in five sub-periods	82
4.A.4	Geometric performance measures for equal growth investors in five sub-periods	82

List of Symbols and Abbreviations

α	Modified standard deviation
γ	Risk aversion parameter
η	Kelly fraction in fractional Kelly strategies
Λ	Strategy
μ	Mean return on the risky asset
μ_p	Mean return on the investor portfolio
$\mu_{p,g}$	Geometric mean return on the investor portfolio
$\hat{\mu}$	Estimator for mean return on the risky asset
σ	Standard deviation of risky asset returns
σ_p	Standard deviation of investor portfolio returns
$\sigma_{p,g}$	Geometric standard deviation of investor portfolio returns
$\hat{\sigma}$	Estimator for standard deviation of risky asset returns
τ	Estimation period length
ϕ	Distribution width parameter
$\psi()$	Return distribution density function
ω	Weight of risky asset in investor portfolio
ω^{MV}	Weight of risky asset in mean-variance investor portfolio
ω^{GO}	Weight of risky asset in growth optimal investor portfolio
∞	Infinity
A	Minimum return on the risky asset
argmax	Argument of the maximum
avg	Average
B	Maximum return on the risky asset
CAPM	Capital Asset Pricing Model
E	Expectation operator
exp	Exponential function

e.g.	Exempli gratia; for example
et. al.	Et alii; and others
f	Fraction of wealth
G	Growth rate of wealth
GO	Growth optimal investor
h	Parameter to adjust return distribution
i.e.	Id est; that is
L	Number of rounds lost
lim	Limit of a sequence or of a function
log	natural logarithm
MV	Mean-variance investor
MV lim	Mean-variance investor who never exceeds growth optimal risk taking
n	Round index
N	Number of rounds
$\mathcal{N}()$	Normal distribution
p	Probability of success
q	Probability of failure
r	Return on the risky asset
r_f	Return on the risk free asset
r_p	Return on the investor portfolio
S	Number of rounds won (success)
SR	Sharpe Ratio
std. dev.	Standard deviation
t	Period index
T	Number of periods
U	Utility function
U.S.	United States
W	Wealth

1 Introduction

In his seminal publication to the field of growth optimal strategies, Kelly (1956) has outlined how a receiver of a noisy signal with information on the outcome of a game can use that information to his advantage in gambling. The work of John L. Kelly was inspired by his Bell Laboratories colleague Claude Shannon and his pioneering work on information theory (Shannon, 1948). Shannon discovered a formula for the maximum rate of transmission of a noisy channel which Kelly applied to gambling¹. This being the first application of that formula to an investment situation, gamblers and investors alike subsequently used what became known as the Kelly criterion for investing and risk management.

Breiman (1961) provided the rigorous mathematical proves of the findings that Kelly (1956) presented. Breiman showed that the use of the Kelly criterion to determine risk taking entails logarithmic utility. While logarithmic utility had already been proposed by Daniel Bernoulli in 1738 (Bernoulli, 1954)², much of 20th century academic interest was triggered by the publications of Kelly (1956) and Breiman (1961) who described that the maximization of logarithmic utility in wealth maximizes the long-run growth rate of wealth. While economists have questioned growth optimal strategies and the underlying assumption of logarithmic utility (Samuelson, 1979), a wealth of successful applications and research in support of it has been developed (Ziemba, 2015).

Portfolio theory as presented by Markowitz (1952, 1959) has become the cornerstone of portfolio selection and the common basic description of an investor is that of a mean-variance optimizer. In this framework, a broad set of concave utility functions is permissible. While its compatibility with a wealth of utility functions makes mean-variance analysis attractive, it essentially is a single-period concept and its generalization to multi-period problems is not straightforward. Optimal multi-period solutions to mean-variance problems for arbitrary utility functions are clearly desirable, but difficult to obtain. For advocates of the pure theory of expected utility maximization it is hard to accept this difficulty as an argument to restrict utility functions, particularly if only logarithmic utility is used. Nevertheless, the

¹See Poundstone (2005) for an enlightening and entertaining discussion of the history of gambling and the Kelly criterion.

²Bernoulli (1954), published in *Econometrica*, is a translation and reprint of the original latin language article published 1738 in *Papers of the Imperial Academy of Sciences in Petersburg V*.

difficulty or even unavailability of theoretically optimal solutions to multi-period mean-variance problems without restrictive assumptions on utility functions and return distributions is certainly a factor that encouraged research on growth optimal strategies.

In this dissertation we focus on comparing growth optimal strategies and myopic mean-variance optimization, which uses the myopic assumption also in settings where myopia does not yield theoretically optimal solutions. We contribute to the existing work on mean-variance optimization by investigating the loss in wealth and the deterioration in performance measures that the myopic assumption can cause. We compare the myopic mean-variance investor with a growth optimal Kelly investor and a fractional Kelly investor to illustrate the lower wealth accumulation and performance. The inspiration to this research goes partly back to Markowitz (1976) who outlines that there is a point on the single-period efficient frontier beyond which one finds portfolios that exhibit greater volatility in the short run, but less return in the long run. This point is the long run growth optimal portfolio and all portfolios that entail more risk taking than the growth optimal portfolio are dominated by the growth optimal portfolio as the time horizon converges to infinity.

First, this dissertation presents a review of theory and applications of growth optimal strategies and the Kelly criterion to provide fundamental knowledge and an overview of what has been done in previous research. Second, this dissertation presents a simulation study on overbetting of myopic mean-variance investors. Overbetting refers to risk taking beyond growth optimal, which induces increased short term volatility and reduced long term growth. We find that a constraint which requires the myopic mean-variance investor to not exceed growth optimal risk taking increases investor performance and prevents bankruptcy in discrete time multi-period asset allocation problems with normally distributed returns. Third, this dissertation presents results for 80 years of stock market investing of a growth optimal investor and a myopic mean-variance investor. Thereby we extend the existing simulation-based work on growth optimal strategies in two ways. First, we use real-world data and have our investor allocate his assets at each decision point given some estimates for next period return obtained as moving averages from past data only. Second, we compare the performance of our growth optimal investor with that of a myopic mean-variance investor who bases his asset allocation decisions on the same parameter estimates. We find that both, the growth rate of wealth and the performance measures achieved by our growth optimal investor substantially surpass those of our myopic mean-variance investors.

2 Growth Optimal Strategies and the Kelly Criterion: A Review of Theory and Applications¹

2.1 Introduction

Growth optimal strategies refer to investment policies which maximize the expected log growth rate of wealth. In the long run, such policies asymptotically outperform any other investment policy in many respects. Applied to an infinite sequence of investments, a growth optimal strategy accumulates asymptotically more wealth than any different investment strategy. Applied to a finite sequence of investments, a growth optimal strategy maximizes the expected growth rate of wealth and never risks ruin. It tends to accumulate wealth quickly, but it exhibits substantial risk taking and thus high volatility in the wealth path. From an expected utility theory viewpoint, growth optimal strategies maximize logarithmic utility. In order to apply a growth optimal strategy in practice, the underlying investment opportunity must be favorable, meaning the expected return must be positive and the maximum losses must be finite.

Modern research on growth optimal strategies originated with the independent papers of Kelly (1956) and Latané (1959), which both suggest strategies to maximize the expected long run growth rate of wealth in multi-period investment situations, which is the compound return. With reference to Kelly (1956), the growth optimal wager became also known as the Kelly criterion. Since Kelly and Latané, growth optimal strategies have been applied in various fields in finance. Growth optimal strategies and the Kelly criterion have predominantly been used in asset pricing, risk management and capital growth problems in investing and gambling.

When portfolio theory as established by Markowitz (1952, 1959) was not yet widely accepted as the premier model of financial markets, growth optimal strategies and logarithmic utility models have been considered a viable alternative to portfolio theory. While portfolio theory comes with restrictive assumptions on the probability distribution of returns, logarithmic utility models instead restrict utility functions. However, when the acceptance of the expected utility maxim by von Neumann

¹This chapter is based on the homonymous working paper Winselmann (2017a).

and Morgenstern (1944) as the criterion for rational choice among risky ventures became more widespread, the initial excitement of the finance community about logarithmic utility models for financial markets decreased. The capital asset pricing model of Sharpe (1964) and Lintner (1965) gained widespread acceptance in the finance community due to its compatibility with a broad range of utility functions, while logarithmic utility models were deemed an inadequate foundation for asset pricing. Despite this demise and some initial imprecise theoretical treatment of their asymptotic properties when time horizons are finite (Samuelson, 1969, 1979), growth optimal strategies are successfully used since half a century by practitioners and academics alike.

Over time and especially in earlier literature, different authors used different terms to refer to growth optimal strategies. In recent publications, the terms *Kelly criterion* (Rotando and Thorp, 1992; MacLean et al., 2010b; Luo et al., 2014) as well as *Kelly strategies* seem to prevail (MacLean et al., 1992; Thorp, 2006; MacLean et al., 2010c; Davis and Lleo, 2013). Further terms that are used to refer to growth optimal strategies and the Kelly criterion include *capital growth model* or *capital growth theory* (Hakansson, 1971a; Bicksler and Thorp, 1973; Hakansson and Ziemba, 1995; MacLean and Ziemba, 2006), *maximum geometric mean strategy* and *average-compound-return criterion* (Samuelson, 1971), *growth optimum model* (Roll, 1973; Kraus and Litzenberger, 1975), *maximum-expected-log (MEL) rule* (Markowitz, 1976), *geometric mean criterion* (Ophir, 1978; Latané, 1978), and *fortunes formula* (Poundstone, 2005; MacLean et al., 2011). The growth optimal portfolio is also frequently referred to as the *numeraire portfolio* in asset pricing applications (Long, 1990; Bajoux-Besnainou and Portait, 1997).

In this paper we wish to add to the debate and systematically review existing research and applications. This paper is organized as follows. First, we explain the Kelly criterion and review the basic literature on the Kelly criterion. This includes a discussion of the expected utility foundations of the Kelly criterion and the critique brought forward against it. Second, we review applications and discuss additions. This covers the fields of portfolio choice and asset pricing, risk and risk management, as well as investing and gambling. We conclude with critical remarks and directions for future research.

2.2 The Kelly criterion

In his seminal publication to the field of growth optimal strategies, Kelly (1956) applies the information theory work of Shannon (1948) to a gambling situation.

Kelly outlines how a receiver of an uncertain private information, which provides him with an edge over the bookmaker, can use this information for profitable gambling. He formulates a strategy which maximizes long-run wealth growth for an infinite series of favorable games.

Kelly's inspiration for the research question came from a widespread fascination in gambling². At the time of Kelly's publication, sports betting on distant events was offered by bookmakers all over the United States, despite being illegal in most states. The bookmakers were dependent on wire services that provided them with timely information on the end and the results of games. If customers were able to obtain information on the outcome of a sports event earlier than the bookmakers, they could make arbitrage profits at the expense of the bookmaker by placing bets on contestants who already won.

Kelly (1956) assumes a gambler who receives some noisy private information on the outcome of a gamble before the bookmaker. Kelly's gambler is able to place a bet given the private information he receives before the bookmaker is informed that results are available and thus closes betting. If the private information was entirely reliable, the optimal strategy of the gambler would be to bet as much as he can. But information transmission at the time of Kelly was slow and blurry, meaning that information transmitted under time pressure is uncertain. In accordance with the information theory developed by Shannon (1948), Kelly assumes that the private information of the gambler is transmitted over a noisy channel. According to information theory, information received through a noisy channel does not have an unambiguous meaning, but it rather has a probabilistic meaning and the sum over the probability of all possible meanings equals one. The receiver of some uncertain information on the results of a sports event has an edge over a yet uninformed bookmaker. Although he cannot be entirely sure what the result is, he can update his probability beliefs and bet accordingly. Kelly (1956) solves the problem of maximizing the long run growth rate of wealth given uncertain private information. Based on the information theory of Shannon, Kelly derives a formula that determines the fraction of wealth to be bet in order to maximize long-run wealth growth in an unending sequence of favorable investments. This formula later became known as the Kelly criterion. Applied to finite sequences of investments, the investment strategy proposed by Kelly maximizes the expected growth rate of wealth. The noisy channel assumed by Kelly is a specific example of a source of an

²See Poundstone (2005) for an enlightening and entertaining discussion of the history of professional gambling and the link of gambling and the Kelly criterion.

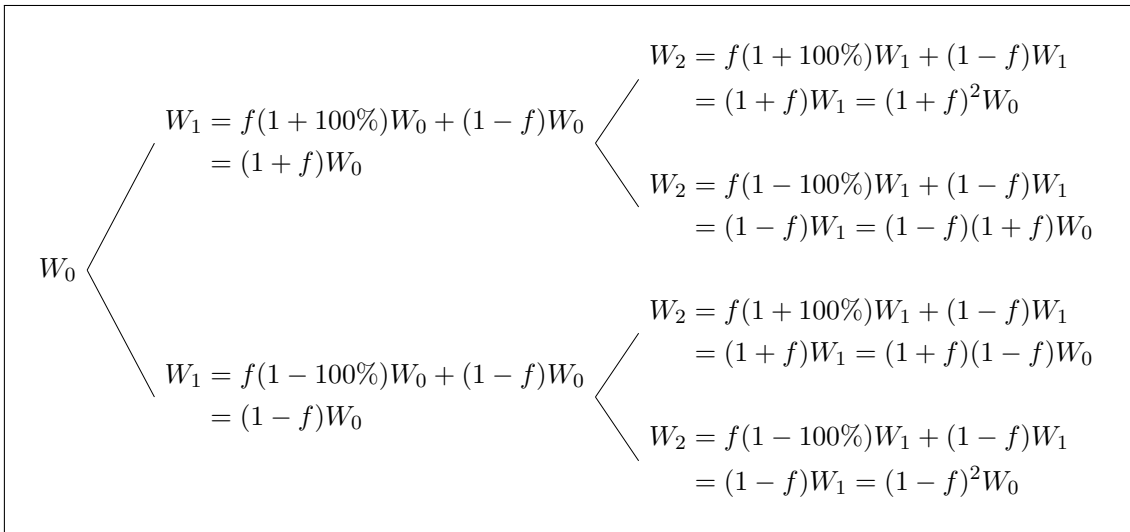


Figure 2.1: Binomial tree depicting the wealth development for two sequential Bernoulli trials where each success returns 100% on the invested fraction f of wealth and each failure results in a loss of the invested fraction of wealth

edge in gambling, but the mathematics developed by Kelly apply in the same way to any favorable investment opportunity.

2.2.1 Computation of the growth optimal wager

The investment problem that Kelly (1956) presented concerns a gamble with a binary outcome. Assume there are two possible outcomes, either success or failure. A success results in a return of +100% and a failure results in a return of -100% on the investment, meaning the investment situation is a double-or-nothing game. Such games are referred to as Bernoulli trials. Let S denote the number of successes and let L denote the number of failures in N sequential investments with $N = S + L$. The investor wealth W_N after N investments is then given by

$$W_N = (1 + f)^S(1 - f)^L W_0 \tag{2.1}$$

where f denotes the constant fraction of wealth invested in each round and W_0 denotes initial wealth. The binomial tree in figure 2.1 visualizes the wealth process

for the first two sequential investments. Kelly (1956) is credited with the idea to maximize the long run log growth rate of wealth G given by³

$$G = \lim_{N \rightarrow \infty} \frac{1}{N} \log \frac{W_N}{W_0} \quad (2.2)$$

Substituting into equation 2.2 the formulation for W_N/W_0 that we obtain from dividing equation 2.1 by W_0 , we obtain

$$G = \lim_{N \rightarrow \infty} \left[\frac{S}{N} \log(1 + f) + \frac{L}{N} \log(1 - f) \right] \quad (2.3)$$

Let p now be the probability of a success and $q = 1 - p$ be the probability of a failure, so that as $N \rightarrow \infty$ we have $p = \frac{S}{N}$ and $q = \frac{L}{N}$. Thus equation 2.3 becomes

$$G = p \log(1 + f) + (1 - p) \log(1 - f) \quad (2.4)$$

and setting the derivative of G with respect to f equal to zero we find

$$\begin{aligned} G' &= \frac{p}{1 + f} - \frac{1 - p}{1 - f} \\ &= \frac{2p - f - 1}{(1 + f)(1 - f)} \stackrel{!}{=} 0 \end{aligned} \quad (2.5)$$

from which follows that the long run growth optimal fraction of wealth to wager in each trial f^* is given by

$$\begin{aligned} f^* &= 2p - 1 \\ &= p - q \end{aligned} \quad (2.6)$$

The growth optimal fraction of wealth to wager became known as the Kelly criterion. It can be generalized to games with uneven payoff where D units are won for every unit invested. Thorp (1984) shows that for such investment situations with a positive expectation, meaning $p > 0$ and $pD - q > 0$, the growth optimal wager is $f^* = \frac{Dp - q}{D}$.

Thorp (1969, 2006) discusses the use of the Kelly criterion in gambles with continuous distributions of outcomes. Rotando and Thorp (1992) provide a generalization of the Kelly criterion to continuous distributions and a stock market example in

³Kelly (1956) shows that the maximum attainable G is equal to the transmission rate of a noisy channel defined by Shannon (1948), which means that a gambler will be able to generate return equal to the rate of transmission of the channel providing him with private information. The analogy to Shannon's information theory was of inspiration and principle concern to Kelly, but his discovery of a formula for the growth optimal level of risk taking and the implied logarithmic utility is today of much more interest to research than the equality to rate of transmission.

discrete time with normally distributed returns. Compared to later research which frequently combines growth optimal strategies with rather complex risk management approaches involving multiple assets and heuristics, Rotando and Thorp (1992) provide an intuitive and accessible introduction to the mathematics of growth optimal wagers.

2.2.2 Expected utility theory foundations

While Kelly (1956) maximizes the long term growth rate of wealth, economists rather resort to utility functions to model the risk preferences of investors. From expected utility theory we know that expected utility is the correct criterion for rational investors to choose among risky alternatives (von Neumann and Morgenstern, 1944). A rational investor would not necessarily maximize wealth, as Kelly (1956) suggests, but he would maximize expected utility. The investment behavior in gambling is only one specific display of the risk preferences of a rational investor, and his risk preferences in gambling should equal his risk preferences in any other investment situation with uncertainty.

Breiman (1961) provides the mathematical foundations to the work of Kelly (1956) and thereby reconciles it with expected utility theory. Breiman rigorously proves that the utility function underlying the strategy proposed by Kelly is log utility in final wealth. An investor with logarithmic utility in final wealth pursues exactly the strategy that Kelly (1956) finds to result in maximum long run capital growth, namely the myopic maximization of the single period expected log growth rate of wealth. The expected utility function that a Kelly investor maximizes is

$$EU = E [\log(W_N)] \tag{2.7}$$

By the properties of the logarithm function, it can be shown that a Kelly investor does not need to solve a complex multi-period problem, but that an optimal policy for a Kelly investor is myopic. Let r_n be the portfolio return in period n .

$$\begin{aligned} E [\log(W_N)] &= E \left[\log \left(W_0 * \prod_{n=1}^N (1 + r_n) \right) \right] \\ &= \log(W_0) + \sum_{n=1}^N E [\log(1 + r_n)] \end{aligned} \tag{2.8}$$

Equation 2.8 shows that logarithmic utility induces myopic investment behavior and that the optimal investment policy is to maximize the expected log growth rate

of wealth in each period independent of all other periods. When the portfolio to be held in period n is chosen, W_{n-1} is known and $\log(W_{n-1})$ is a constant as far as this choice is concerned (Markowitz, 2014).

Breiman (1961) does not only provide the rigorous proof that logarithmic utility is the underlying utility function of the growth optimal strategy proposed by Kelly (1956), but Breiman also proves that an investor who maximizes expected logarithmic utility in final wealth asymptotically accumulates more capital than any other investor. In more technical terms, Breiman (1961) establishes that independent of the distribution of returns, the final wealth W_N^* accumulated under a growth optimal strategy Λ^* will asymptotically be infinitely larger than the final wealth W_N accumulated under any essentially different strategy Λ , which is not asymptotically close to Λ^* . If a non-terminating growth optimal strategy Λ^* achieves wealth W_N^* after N plays and a strategy Λ achieves wealth W_N , then there exists a limit for W_N^*/W_N and $\lim_{N \rightarrow \infty} (W_N^*/W_N) = \infty$. The optimal strategy Λ^* , which is the maximization of the expected logarithm of final wealth, asymptotically dominates any other essentially different strategy, both in terms of the final wealth and the time required to reach some distant wealth goal.

Before Breiman (1961), Latané (1959) introduces logarithmic utility to finance independent of Kelly (1956) and argues that it has superior long-run properties, meaning that it ensures highest long run wealth growth. Neither of these authors has though been the first to imply logarithmic utility. The idea that there must be a diminishing marginal utility of wealth and that a logarithmic function might be appropriate to describe it can be traced back as far as to Daniel Bernoulli in 1738 (Bernoulli, 1954)⁴, who describes a logarithmic utility function and applies it to solve the St. Petersburg paradox. Although Kelly and Latané were not the first scholars to propose logarithmic utility, their work is recognized as an important step to apply logarithmic utility to problems in finance and bring it back into modern academic discussion.

Investor preferences are in later research rather represented more broadly by the family of isoelastic utility functions, of which logarithmic utility is a special case. The family of isoelastic utility functions is defined as

$$U = \begin{cases} \frac{W^{1-\gamma}-1}{1-\gamma} & \forall \gamma \neq 1 \\ \log(W) & \forall \gamma = 1 \end{cases} \quad (2.9)$$

⁴Bernoulli (1954), published in *Econometrica*, is a translation and reprint of the original latin language article published 1738 in *Papers of the Imperial Academy of Sciences in Petersburg V*.

where γ denotes risk aversion. $\gamma = 1$ represents the special case of logarithmic utility, which can be established by virtue of l'Hôpital's rule. Mossin (1968) discusses multi-period portfolio policies and establishes the conditions under which myopic portfolio choice is optimal. Mossin finds that with isoelastic utility, the investor preferences are such that a stationary policy and complete myopia are optimal. Isoelastic utility functions exhibit the property of constant relative risk aversion, which is a necessary and sufficient condition for myopia being optimal (Mossin, 1968). Hakansson (1971b) confirms the results of Mossin (1968) for logarithmic utility, but clarifies that myopia is only optimal for isoelastic utility in general if independence of period by period returns is assumed. Logarithmic utility induces complete myopia even when yields over periods are nonindependent.

The results of Mossin are used by Samuelson (1969), who investigates lifetime portfolio selection by dynamic stochastic programming in discrete time. Samuelson finds that isoelastic utility and independence of returns over time splits the Ramsey-Phelps savings problem (Phelps, 1962) and the portfolio selection problem into two distinct problems. The decision which fraction of wealth to consume and which fraction of wealth to invest becomes independent of time and the level of wealth, meaning that consumption and investment will be constant fractions of wealth in each period. Also, the portfolio decision of how much to invest into a risky asset versus a risk free asset becomes independent of consumption, time and the level of wealth. Samuelson (1969) thereby rebuts the idea that risk tolerance decreases towards an investor's end of life and proves this for all isoelastic utility functions. Merton (1969, 1971) presents equivalent findings for the continuous time setting.

Logarithmic utility and—if independence of returns over time is assumed—isoelastic utility thus render optimal portfolio policies myopic, which is a desirable property for practical applications. For any other set of assumptions and utility functions, myopia in multi-period problems does not lead to optimal outcomes. Multi-period problems that cannot be treated as myopic have to be either solved recursively or heuristically. Recursive approaches to multi-period portfolio problems are complex and computationally intensive. In addition, they frequently require assumptions on numerous future parameters⁵ that are yet unknown today and difficult—if not impossible—to estimate (Kolm et al., 2014).

A differently styled defense of logarithmic utility is provided by Sinn (2003). The argument is that in the biological selection process, expected utility maximization

⁵See Chopra and Ziemba (1993) for an investigation into the effects of errors in parameter estimates on portfolio choice.

under a logarithmic preference function is dominant in the sense that a population with any other risk preference will disappear relative to a population with logarithmic preferences for economic decisions under uncertainty. In the limit, as time goes to infinity, this will occur with a probability that approaches certainty. A similar argument is made by Amir et al. (2005) in the context of financial markets.

2.2.3 Fractional Kelly strategies

The risk taking suggested by growth optimal strategies is often perceived as high, meaning that the wealth path is somewhat rough and the volatility of wealth levels over time is high. In order to reduce risk, fractional Kelly strategies have been suggested. Fractional Kelly refers to investing a specific fraction of the full Kelly growth optimal wager, such as one half or one quarter, in order to increase security at the expense of growth. MacLean et al. (1992) credit Friedmann (1982) to have been the first to use fractional Kelly strategies in an unpublished manuscript, although already Breiman (1961) and Thorp (1969) mentioned that systematically lower than growth optimal risk taking in repeated gambles results in a reduced expected growth rate of wealth and less volatility in the wealth path.

The justification of such strategies is not as straightforward as that of full Kelly strategies. MacLean and Ziemba (1999) argue that the theoretical justification of fractional Kelly strategies can be made using a continuous time approach with the assumption of lognormal assets. MacLean et al. (2005) further point out that for lognormal investments η -fractional Kelly wagers correspond to using power utility, such as the isoelastic utility function presented in equation 2.9 above, where $\eta = \frac{1}{\gamma}$. Ziemba (2015) argues that this result can also serve as an approximation for non-lognormal assets. Davis and Lleo (2013) review the literature on fractional Kelly strategies in continuous time. MacLean et al. (1992) refer to fractional Kelly strategies in discrete time settings with normally distributed returns as an effective trade-off of growth versus security, although not fully efficient. MacLean et al. emphasize their ease of computation as a principal argument for the use of fractional strategies in investment situations where they are theoretically suboptimal.

Davis and Lleo (2010) define fractional Kelly strategies differently in order to obtain optimal results also for non-lognormal asset. Davis and Lleo build on their earlier work (Davis and Lleo, 2008) and extend Kelly and fractional Kelly strategies to benchmark oriented asset management as well as asset and liability management. In the benchmarked case, the spectrum of chosen portfolios lies between the replication of the benchmark and the full Kelly growth optimal portfolio for η -fractional Kelly strategies with $0 \leq \eta \leq 1$. Davis and Lleo (2013) investigate the conditions

that ensure optimality of fractional Kelly strategies when assets are non-lognormal. Davis and Lleo find that optimality of fractional Kelly strategies can be ensured by having the definition of fractional Kelly strategies coincide with the fund separation theorem of the underlying problem.

2.2.4 Critique of the Kelly criterion and logarithmic utility

We have outlined above that the utility foundation of the Kelly criterion and growth optimal strategies is logarithmic utility. For any investor whose utility function is not logarithmic, it is not utility maximizing to pursue a full Kelly strategy, although the wealth accumulation of this strategy asymptotically dominates that of any other strategy in the long run. The statement that full Kelly is only optimal if utility is logarithmic may seem obvious, but it has been at the center of a heated discussion in numerous publications and so were also the asymptotic properties of growth optimal strategies. This section provides an overview of critique brought forward against the initial imprecise treatment of the asymptotic properties of growth optimal strategies and the attempted generalization of logarithmic utility.

Some early advocates of logarithmic utility based models of financial markets have either implied that the utility function of all investors can be approximated by logarithmic utility, or sidelined utility theory in general. Rubinstein (1976) explicitly states a preference for models that restrict utility functions over models that restrict return distributions. Latané (1959) is criticized for advocating growth optimal strategies as an alternative to expected utility maximization (Samuelson, 1969; Ophir, 1978). These views encountered fierce opposition from proponents of expected utility theory. Samuelson opposed growth optimal strategies on grounds of the view that individuals have different risk preferences and that the approximation of all individuals' risk preferences by logarithmic utility is poor, regardless of whether or not it has computational advantages or desirable asymptotic properties in the long run (Samuelson, 1969, 1971, 1979; Merton and Samuelson, 1974). Some mathematical deficiencies of early papers on growth optimal strategies and the theoretical arguments against the generalization of logarithmic utility brought forward by Samuelson are fully accepted by proponents of growth optimal strategies today (MacLean et al., 2010b), but the partly implied conclusion that growth optimal strategies are generally without merit is rejected. MacLean et al. (2010b) and Ziemba (2015) suggest that because Samuelson objected to growth optimal strategies and pointed at deficiencies in selected authors works, the field of growth optimal strategies in general has received less attention by researchers than it deserves.

Notably Latané (1959), but also some other proponents of the maximization of the geometric mean have been criticized for a too loose application of the law of large numbers or the central limit theorem (Samuelson, 1971). Indeed, a growth optimal strategy asymptotically dominates the wealth accumulation of any essentially different strategy only as the length of the sequence of re-investments approaches infinity. Careful attention has to be paid to the asymptotic nature of this property. As Samuelson (1971) points out, this does not imply that growth optimal strategies are optimal for any finite time horizon. The only case where growth optimal strategies are indeed optimal for finite time horizons is, from an expected utility point of view, the case of logarithmic utility—and in this case the reasoning is quite different. Growth optimal strategies are not optimal for logarithmic utility because they are growth optimal, but because they maximize expected utility (Samuelson, 1979). A series of published replies and counter-replies to Latané (1959) by Ophir (1978, 1979) and Latané (1978, 1979) has centered on a related issue. As correctly pointed out in these replies, growth optimal strategies do not maximize the long run growth rate, but the mathematical expectation of the long run growth rate. They do not maximize terminal wealth, but in the long run they lead to an almost certainly higher terminal wealth than any different strategy. The favorable properties of growth optimal strategies are asymptotic, and for them to materialize with certainty the length of the time horizons has to converge to infinity.

As Rubinstein (2006) puts it, the asymptotic properties of growth optimal strategies appear so attractive that some researchers have falsely conjectured that maximization of the expected logarithm of final wealth and thus the geometric mean is the best long-run strategy for any investor. At least since Mossin (1968) we know that this conjecture is mistaken. The maximization of the expectation of the geometric mean return is only optimal if utility is logarithmic. Proponents of logarithmic utility praise it for its analytic and computational ease as well as for its superior performance in some practical applications and for its lower sensitivity to estimation errors than portfolio theory. These aspects have to be traded off against the theoretical arguments of Samuelson, as the maximization of the expectation of the geometric mean return either requires an investor who pursues some objective outside expected utility theory, or it requires the restrictive assumption of logarithmic utility. Fama and MacBeth (1974) reach mixed conclusions on whether or not the stock market is from an empirical perspective dominated by growth optimal investors. Fama and MacBeth tend to reject the hypothesis that logarithmic utility is the prevalent utility function in the price formation process. Notwithstanding, growth optimal strategies find numerous applications in finance.

2.3 Applications

Based on the research findings outlined in the previous section, applications of growth optimal strategies have emerged. In this section we review applications presented in the literature from the inception of modern research on growth optimal strategies starting with Kelly (1956) until today. We cluster research in three broad areas. These are portfolio choice & asset pricing, risk management, and investing. Portfolio choice and asset pricing refers firstly to the initial attempts to develop capital market models that build on logarithmic utility of investors and secondly to more recent approaches which use the growth optimal portfolio as a numeraire portfolio to simplify complex asset pricing problems. Risk management refers to research that discusses the risk in the form of volatility in the wealth path and final wealth that growth optimal strategies exhibit, as well as to the avoidance of investor ruin that a maximization of the expected geometric growth rate entails. Investing refers to more practical applications in gambling and investing alike. Gambling and investing are essentially equivalent from a mathematical point of view. The exploitation of card-counting strategies in black jack is equivalent to the exploitation of suspected stock market anomalies.

2.3.1 Portfolio choice and asset pricing

At the outset of portfolio theory (Markowitz, 1952, 1959), some academics considered it a deficiency of the model that it did not suggest a specific portfolio to an investor, but only gave him a set of efficient portfolios among which he ought to decide himself in order to maximize his expected utility (Roy, 1952; Latané, 1959). In addition, Markowitz' portfolio theory was a single period concept that proved to be difficult to extend to multi-period time horizons that more accurately reflect the real life problem of investors (Markowitz, 1959, 1976). The interested reader is referred to Steinbach (2001) for a more encompassing review of mean-variance models.

Samuelson (1969) analyses discrete time formulations and Merton (1969, 1971) analyses continuous time formulations of portfolio choice problems. Despite constituting important advancements, both do not provide a solution for discrete time problems when utility functions are not of the isoelastic type. Also, Samuelson (1969) requires the assumption of independence of returns over time, which faces severe challenges in real-world applications. Both, the fact that logarithmic utility based models recommend a specific asset allocation to investors (Latané, 1959) and the fact that logarithmic utility models suggest investment decisions that are optimal for multi-period time horizons without further restrictive assumptions (Breiman,

1961) added to their appeal. These perceived advantages were motivation to the attempts to use logarithmic utility based models as premier models of financial markets. Hakansson argues in favor of the logarithmic utility based growth optimum model for two reasons. Firstly, Hakansson (1971b) investigates the properties of growth optimal strategies—such as consistency with expected utility theory if log utility is assumed, decreasing absolute risk aversion with respect to variance, and investor survival—which he finds to be well suited for a portfolio model. Secondly, Hakansson (1971c) remarks that the myopic property of the logarithmic utility function is desirable given the difficulty to estimate future yields and correlations. Hakansson (1971a) reemphasizes these arguments and presents a simulation study of wealth accumulation in the intermediate run. After only six periods, he finds the growth optimal strategy to be highly likely to outperform other strategies, while being less likely to encounter investor ruin. Based on the entirety of his findings, Hakansson considers the mean-variance model to be severely compromised by the capital growth model.

Roll (1973) discusses evidence on the growth optimum model. Starting with the line of arguments presented by Hakansson (1971a), Roll emphasizes the investor ruin that a sequence of re-investments in single-period mean-variance efficient portfolios can bring. From today's perspective, knowing that investor ruin in multi-period mean-variance problems frequently stems from incorrect applications of the single-period mean-variance framework to multi-period horizons, the publication by Roll highlights the challenges that the extension of mean-variance analysis to multi-period problems brought at that time. Further, Roll (1973) finds in an empirical investigation that the results of mean-variance analysis and the growth optimum model can be close to indistinguishable when the utility function used in mean-variance analysis is a quadratic approximation to the logarithm function. Kraus and Litzenberger (1975) confirm the findings of Roll (1973) in a more general setting. Bicksler and Thorp (1973) compare growth optimal strategies with State Preference Theory (Arrow and Debreu, 1954) and the Capital Asset Pricing Model (Sharpe, 1964; Lintner, 1965). Bicksler and Thorp build on Roll (1973) and support his findings with further numerical evidence obtained from simulation results for the distribution of final wealth over different time horizons.

Fama and MacBeth (1974) emphasize that while the growth optimum model and mean-variance analysis were often presented as two distinct and theoretically incompatible approaches, they actually are not. Fama and MacBeth (1974) empirically investigate whether the theoretical properties of growth optimal portfolios are consistent with the properties of observed portfolios that real world investors actually

hold. They test two hypotheses. The first hypothesis is that the price-formation process is dominated by growth optimizers, which they cannot reject. The second hypothesis is that the market portfolio is the growth optimal portfolio, for which Fama and MacBeth find mixed evidence. While the hypothesis cannot be rejected on the basis of standard significance tests, a detailed examination of the data and comparison of portfolios does suggest structural differences, with the market portfolio being systematically less risky than the growth optimal portfolio. Fama and MacBeth interpret their results as evidence against Roll's (1973) interpretation of his empirical findings.

Rubinstein (1976) calls for logarithmic utility based models to become the premier models of financial markets. Most models require simplifying assumptions. The simplifying assumption of logarithmic utility models is that utility is restricted to be logarithmic. Portfolio theory usually is used in connection with restrictive assumptions concerning the probability distribution of returns. Rubinstein (1976) argues that models with restrictive assumptions on utility functions are superior to those with restrictive assumptions on probability distributions.

Markowitz (1976) takes a decisive step to reconcile the literature on the growth optimal model with mean-variance analysis. By then, mean-variance analysis had become more widely accepted and so was the notion that investor preferences ought to be modeled by utility functions. Markowitz (1976) emphasizes that the growth optimal portfolio choice is one of the possible portfolio choices along the single-period efficient frontier and all portfolio choices that take more risk than the growth optimal portfolio are dominated by the growth optimal portfolio when the time horizon converges to infinity. Markowitz welcomes the insights of growth optimal strategies and considers them to add value to his theory of portfolio selection by indicating the maximum level of risk that a "conservative investor" (Markowitz, 1976) should take if the time horizon is long.

The concept of stochastic dominance has been the basis of a further attempt to create a theory of security prices. Jean (1980) establishes a link between stochastic dominance and growth optimal strategies. Similarities between growth optimal portfolios and stochastic dominance portfolios are reviewed and connected to earlier findings of Hakansson (1971a,b). Jean (1980) proposes to use the geometric mean maximization algorithms which were developed in the field of growth optimal strategies to further explore stochastic dominance portfolios. Jean finds that a geometric mean ranking is a necessary condition for any stochastic dominance ranking of portfolios, albeit not a sufficient condition for stochastic dominance.

Academic work on using the growth optimal portfolio in asset pricing has since then continued. Notable efforts include the use of a numeraire portfolio in derivatives pricing, risk management, portfolio optimization (Platen, 2006). Long (1990) introduces the numeraire portfolio as a portfolio relative to which all other assets can be priced. The numeraire portfolio is equivalent to the growth optimal portfolio for a given list of assets. Long (1990) distinguishes the numeraire portfolio from the growth optimal portfolio as the application he intends is to simplify the representation of asset prices by expressing prices relative to the numeraire portfolio, which may be done for a limited set of assets in which case the numeraire portfolio is not growth optimal with respect to the market at large. Bajeux-Besnainou and Portait (1997) review the literature on the numeraire portfolio, the properties of the portfolio, and the extensive links it has to the continuous-time versions of multiple capital market theories. Platen (2006) extends the work of Long (1990) and develops a framework for asset pricing which is based on the numeraire approach. Platen (2006) uses the growth optimal portfolio as numeraire portfolio for asset prices and finds this portfolio to be a combination of the market portfolio and the savings account. Du and Platen (2014) use the growth optimal portfolio based benchmark approach of Platen (2006) to manage risk in the form of not perfectly replicable contingent claims. Platen and Rendek (2011) approximate the growth optimal portfolio by naive diversification and show their approximation to perform better and be more stable over time than a sample-based Markowitz mean-variance approach, which is found to be highly sensitive to estimates changing over time. Baldeaux et al. (2015) use the benchmark approach of Platen to price long-dated currency derivatives. Baldeaux et al. find that under their stochastic volatility model, when calibrated to real data, the risk neutral approach fails while the benchmark approach performs better. In the light of this result, they conclude that the long standing practice to adopt the risk neutral pricing paradigm may be unsuitable particularly for long-dated derivatives.

The recent publications on growth optimal strategies and the closely related numeraire portfolio show that the stream of research which uses the growth optimum portfolio in asset pricing and portfolio choice is as vivid as it has been in the early days after the initial publications by Kelly (1956) and Latané (1959). We note that in more recent literature the perspective has changed. Current literature does not present the growth optimum model as theoretically superior to Markowitz' portfolio theory. The attempt is not to overcome limitations of portfolio theory with a new model. Instead, the compatibility of both models and the mean-variance efficiency of the growth optimal portfolio is emphasized. The growth optimal portfolio is either

used as a benchmark risk/return combination, or it is used as an easier to obtain and more stable portfolio than those that can be obtained from portfolio theory under real-world constraints.

2.3.2 Risk and the Kelly criterion

One of the properties of growth optimal strategies is that, while the volatility in final wealth is high, they never risk ruin (Latané, 1959). A higher than growth optimal level of risk taking leads to increased short term variance and reduced expected long-term growth (Breiman, 1961; Markowitz, 1976). Taking twice as much risk as growth optimal is proven to yield high fluctuations of wealth around the initial wealth with an expected long term growth rate of zero (Thorp, 1984).

The relation of risk and the Kelly criterion is ambiguous. On the one side, Kelly strategies are known to be risky. From an expected utility theory point of view, Kelly strategies are only optimal when utility in final wealth is logarithmic. The logarithmic utility function is the low risk aversion $\gamma = 1$ member of the family of isoelastic utility functions (equation 2.9) and it is less risk averse than the majority of individuals is suggested to be (Markowitz et al., 1994). On the other side, the Kelly criterion is known as a risk management tool to prevent excessive risk taking, also known as overbetting, which is seen as a principal cause of hedge fund disasters and gambler's ruin (Thorp, 1984; Ziemba and Ziemba, 2013). Aucamp (1978) correctly remarks that the avoidance of ruin is not a property specific to Kelly strategies, but that any strategy with a geometric mean rate of return ≥ 1 avoids ruin and so does the maximization of the geometric mean in the presence of correct beliefs of the probability distribution of outcomes.

The risk of Kelly strategies in the form of high volatility in final wealth outcomes can be shown in simulation studies. The first published simulation study on the Kelly criterion and growth optimal strategies is Bicksler and Thorp (1973), who simulate terminal wealth for different investor time horizons and leverage. They find that the difference between maximum and minimum terminal wealth outcomes is huge already after short time horizons of 10 to 40 periods and that terminal wealth repeatedly is below initial wealth, but always positive. Bicksler and Thorp confirm that higher leverage does not necessarily lead to higher mean final wealth. Instead, exceeding a certain level of leverage can harm wealth growth, which is in line with theory (Breiman, 1961). Applying the growth optimal concept to realistic stock market return distribution data, Bicksler and Thorp find the growth optimal portfolio to be highly leveraged. Hausch and Ziemba (1985) simulate growth optimal strategies in horse racetrack betting and extend the work of Bicksler and Thorp

(1973) with simulations with longer horizons and with the addition of fractional Kelly strategies. All studies find high growth rates of wealth, but also substantial volatility of terminal wealth outcomes for finite time horizons. MacLean et al. (2011) provide an overview and extension of simulation studies on the medium and long term performance of Kelly strategies. MacLean et al. (2011) assess the performance of growth optimal strategies by determining the probabilities of achieving some wealth goal versus losing some fraction of wealth with full and fractional Kelly wagers.

Rubinstein (1991) takes a more applied perspective and analyses investments into the stock market following a growth optimal strategy. Rubinstein finds that while the Kelly strategy eventually accumulates more wealth than other strategies, it requires long time horizons to outperform other strategies with high probability. Common other strategies investigated in the paper are the all-cash and all-stock strategy, for which Rubinstein finds that the required time horizon to outperform these strategies with 95% certainty is 208 and 4,700 years, respectively. Rubinstein builds on earlier research by Leibowitz and Krasker (1988), who found that for an all-stock portfolio it takes 123 years to outperform an all-bond portfolio with 95% probability, assuming lognormal returns with mean excess return of 2.5% and volatility of 18%. It is worth noting that the rather low excess return assumption considerably adds to the length of the time horizon.

From expected utility theory we know that every rational investor should maximize his expected utility. One might argue that following this route would render growth optimal strategies as a risk management tool obsolete, as all investors only take the risk that maximizes their expected utility. In practice, this is doubtful for at least two reasons. Firstly, it would require investors to act purely rational. But instead, Las Vegas is much rather filled with risk seekers who enjoy the thrill of gambling, than with rational utility maximizers in the sense of von Neumann and Morgenstern (1944). Studies find that also racetrack bettors are risk seekers (Ali, 1977). Secondly, it would require that scientifically exact maximization of expected utility is always possible. Reality is different and uncertainty about input parameters frequently renders exact maximization of expected utility impossible. In line with this, Michaud (1989) claims that mean-variance optimizers are estimation error maximizers. For these reasons, the Kelly criterion is known as a powerful risk management tool in professional gambling and investing (Thorp, 1984; Poundstone, 2005). The growth optimal level of risk taking can be a helpful guidance for any investor or gambler who faces a favorable investment situation. In the presence of parameter uncertainty, a combination of growth optimal strategies with robust optimization approaches which account for possible parameter estimation errors can

provide helpful guidance in risk management. Exceeding the growth optimal level of risk taking may endanger financial survival and reduce long term growth. Ziemba and Ziemba (2013) and Ziemba (2015) assert that Long Term Capital Management and most other hedge fund disasters are rooted in higher than growth optimal risk taking and that the use of the Kelly criterion as a risk management tool could have helped to prevent these disasters.

A practical limitation in the computation of the growth optimal level of risk taking is that it requires the probability distribution of outcomes. Using a flawed approximation or erroneous parameter estimates for the probability distribution of outcomes naturally distorts the result. Correct estimation of tail risk is essential in order to be able to meaningfully use the Kelly criterion to prevent excessive risk taking that decreases long term growth and may endanger financial survival.

2.3.3 Kelly strategies in gambling and investing

Hausch et al. (1981) investigate the efficiency of the market for racetrack betting in horse races. A betting scheme for place and show bets is developed, which builds on the logarithmic utility approach of Kelly strategies. The betting scheme is shown to be profitable, which demonstrates a market inefficiency in the market for race-track betting, contradicting earlier work that suggested there were no inefficiencies large enough to be profitable to exploit (Snyder, 1978). Ziemba and Hausch (1985) extend the work of Hausch et al. (1981) with the addition of fractional Kelly strategies and with the study of the effect of transaction cost in the form of the track take. Transaction costs, such as the track take, make investment opportunities less favorable and the growth optimal level of investment decreases accordingly. Hausch and Ziemba (1990) study arbitrage opportunities if different odds are offered by different bookmakers, in their case racetracks that allow betting on horse races that take place elsewhere. A growth optimal betting model is developed to exploit the arbitrage opportunities that different odds offer.

Clark and Ziemba (1987) investigate possibilities to exploit the turn-of-the-year effect using growth optimal strategies. They measure risk by the probability to reach some multiple of wealth before halving wealth as well as by the probability of reaching some level of wealth before ruin. Note that the possibility of ruin stems from estimation risk. If the probability believes were correct, growth optimal strategies would result in the probability of ruin being zero. Clark and Ziemba (1987) find that during the time period 1977-1987 the turn-of-the-year effect at multiple stock exchanges offered various possibilities to establish profitable trading strategies. All strategies aim at exploiting the historically higher gains that stock markets exhibit

in the month of January. Several publications have since dealt with explaining the turn-of-the-year effect (Ritter, 1988; Lakonishok and Smidt, 1988) and exploiting it (Ziemba, 1994, 2012). An overview of literature on this topic and various other calendar anomalies is provided in Ziemba (2011, 2012).

MacLean et al. (1992) investigate the growth versus security trade-off offered by fractional Kelly strategies. While in their discrete time setting fractional Kelly strategies are not theoretically optimal, MacLean et al. consider them to be effective, meaning that they compensate a loss of expected growth with a gain in security. They measure security by the probability to double wealth before halving it. MacLean et al. (1992) apply fractional Kelly strategies to Blackjack, horse racing, lotteries, and stock market anomalies such as the turn-of-the-year effect. As for any Kelly strategy, it is crucial to have a game with a positive expectation, i.e. to obtain an edge over the casino, bookmaker, or any other counterpart. In Blackjack, the edge over the casino can be obtained from techniques such as card-counting. MacLean et al. state that the edge of a successful card counter varies from -10% to $+10\%$ which he can translate into an average edge of about 2% by wagering more when the card deck is favorable. They find that the probability to double wealth before halving it is 67% for full Kelly and 89% for half Kelly, which provides 75% of the expected growth rate of full Kelly. In horseracing, the edge comes from market anomalies and irrational bettors (Hausch et al., 1981). The market for horseracing follows a particular and complex set of rules, of which Hausch and Ziemba (1995) provide an extensive overview and illustrate favorable betting strategies. In lotto games, MacLean et al. (1992) find the edge to be much smaller and it can mostly be obtained from betting on unpopular numbers. Over all numbers, the expected return per \$1 wagered is only \$0.4 to \$0.5. While the amounts of wins are very high, the probabilities of wins in lotto games are extremely small. Even for the more attractive lotto regimes, such as Canada's 6/49, MacLean et al. find Kelly strategies to be barely attractive to investors since the optimal wagers are fractions of wealth in the range of one dollar per hundred thousand dollar to one million dollar of investor bankroll. While the probability of tenfolding wealth before halving it is attractively high at 95% with a one quarter fractional Kelly strategy, MacLean et al. find that it would take several trillion lottery plays, meaning several billions of years, to achieve this wealth increase. For less favorable lotto regimes in other markets the prospects are even substantially worse, if profitable betting is possible at all. MacLean and Ziemba (1999) and Bain et al. (2006) extend the analysis of MacLean et al. (1992) with further details on specific investment opportunities, such as particular horse races.

MacLean et al. (2004) consider the discrete time problem of investment in risky assets with the objective of achieving maximum capital growth while controlling for downside risk with either value at risk or drawdown constraints. Discrete approximations of the lognormal distributions of asset returns are used. MacLean et al. develop a heuristic which reduces the number of possible outcomes of their three-asset universe to four equally likely scenarios and their proposed algorithm returns an optimal solution to the constrained capital growth problem given the selected scenarios. An additional algorithm to find optimal strategies that trade off growth for security is presented. MacLean et al. (2004) find that for high levels of risk control, the optimal strategy becomes increasingly conservative towards the end of the planning horizon. A growth optimal strategy is optimal for low levels of risk control while investment strategies deviate increasingly from the growth optimal strategy as risk control is tightened. MacLean et al. (2004) find fractional Kelly strategies to be generally not optimal in their discrete time capital growth problem with security constraints.

MacLean et al. (2005) study the performance of fractional Kelly strategies in a continuous time setting with jointly lognormal asset prices with random parameters. Investment risk is represented by a geometric Brownian motion asset price model with random rates of return. The risk that MacLean et al. control for in this study is primarily estimation risk. At each re-balancing time, a Bayes estimator of distribution parameters is defined. The proposed algorithm does not use periodic re-balancing, but rather re-balances when control limits are violated which indicate that asset returns develop inconsistent with their estimates. MacLean et al. (2005) chose to minimize time-to-wealth goals under the restriction of not falling below some wealth threshold with a given probability. The wealth goal approach is compared with a periodic re-balancing expected utility maximization approach and found to be superior, which is mainly attributed to the control-limit induced re-balancing as opposed to re-balancing at fixed time intervals.

MacLean et al. (2006) study the investment of capital in risky assets in a continuous time setting, where asset prices again follow a geometric Brownian motion model with drift and stochastic parameters. The investment problem is composed of two separate problems, one is the optimal growth portfolio and the other is the fraction of wealth to invest in order to control risk. The distribution of future asset returns is estimated from past data and upper and lower boundaries for the wealth process are defined. Re-balancing times are random and re-balancing is induced by violation of the upper or lower boundary of the wealth process, indicating that the trajectory of wealth is not proceeding as anticipated and parameter estimates

require revision. Control limits are determined heuristically. MacLean et al. (2006) find their dynamic investment system to be effective and find it to provide results superior to pure value-at-risk or conditional value-at-risk based risk management. Firstly, this adds to a body of research that is increasingly critical towards value-at-risk based risk management, such as Basak and Shapiro (2001). The proposition of an alternative risk management approach based on growth optimal strategies with control limits as presented in MacLean et al. (2006) is novel. Secondly, as in MacLean et al. (2005), also MacLean et al. (2006) find fixed period re-balancing to yield inferior results compared to control-limit induced rebalancing.

Real-world constraints and risk control render capital growth problems difficult or impossible to solve exactly and optimally. Luo et al. (2013) develop a dynamic portfolio management model based on a simulated annealing algorithm to solve capital growth problems under real world constraints. A drawdown approach to risk management is presented, which is shown to be able to reproduce a variety of common risk measures such as the risk of ruin, variance measures, and shortfall measures. Luo et al. (2014) abstract further from the drawdown constraint and present a more general application of the simulated annealing algorithm for optimal capital growth.

MacLean et al. (2015) study capital growth problems when asset price dynamics are regime-dependent and follow a defined geometric Brownian motion within each regime. While value-at-risk constraints are frequently formulated for end of horizon wealth, this research uses a value-at-risk constraint that has to be met at each point in time with a convex shortfall penalty for violations. This ensures that not only the probability of losses, but also the size of losses is penalized. A deterministic program is developed to calculate optimal wagers at discrete points in time. MacLean et al. (2015) find the convex penalty function to smooth the wealth path and provide an effective trade off between risk and capital growth.

A different stream of research attempts to explain the investment behavior that investors exhibit in practice. The portfolios of endowments and renowned investors frequently exhibit a high concentration in single investments as well as long term returns superior to those of stock markets. Ziemba (2005) claims that the observed portfolios and investment decisions of such investors resemble those of full or fractional Kelly strategies more than any other investment approach. Gergaud and Ziemba (2012) and Ziemba (2015) extend the analysis with additional data and re-confirm the findings of Ziemba (2005) with respect to both, the high concentration in single positions that portfolios of great investors tend to exhibit and the consistently high performance. Great investors in these papers refers to investors such as John

Maynard Keynes, who managed the King's College endowment, or Warren Buffet and Berkshire Hathaway, as well as George Soros and his Quantum fund. Ziemba (2015) finds that the investment behavior and performance of these investors can hardly be explained with anything but Kelly strategies.

2.4 Conclusion

In this paper we review the Kelly criterion and growth optimal strategies. We present a breadth of applications, specifically in the fields of asset pricing, risk management and capital growth problems. We maintain that attempts to use log utility models as premier models of financial markets and as alternative to portfolio theory are efforts of past research. Today, the compatibility of the growth optimum approach with portfolio theory is emphasized instead and the growth optimal portfolio is presented as one possible portfolio on the efficient frontier. What remains of the growth optimal approach in the field of asset pricing is centered around the idea of the numeraire portfolio and the benchmark approach. In risk management, the primary usefulness of growth optimal strategies and the Kelly criterion is to prevent overbetting. While the use of investment policies tied to logarithmic utility has been debated, we shown that multiple successful applications to capital growth problems exist. The computational ease and the lower sensitivity to estimation errors are among the principal advantages of investment policies based on logarithmic utility in final wealth. While growth optimal strategies may not exhibit the full theoretical beauty and comprehensiveness of expected utility theory, it remains important to emphasize both, the feasibility and the partly superior results that these strategies yield under real world constraints.

3 Using the Kelly Criterion to Protect Wealth Growth in Myopic Multi-Period Mean-Variance Portfolio Selection¹

3.1 Introduction

The mean-variance analysis of Markowitz (1952) and his seminal publication on portfolio theory (Markowitz, 1959) provide the fundamental basis for portfolio selection. Although originally cast in a single-period framework, both academia and practice until today widely use mean-variance analysis also in multi-period settings. As true multi-period formulations of mean-variance problems are complex, a common approach is to assume myopia despite the analytic shortcomings of the myopic assumption (Basak and Chabakauri, 2010). While Samuelson (1969) and Merton (1969, 1971) have presented settings in which myopic behavior is optimal, Roll (1973) has pointed to the existence of single-period mean-variance efficient portfolios which bring about investor ruin in an infinite sequence of re-investments. Markowitz (1976) notes that portfolios exist which are from a utility perspective optimal for a single-period horizon, but in the long run cause increased volatility in an investor's wealth path at the expense of long term growth. In this paper we investigate an effect similar to that observed by Roll (1973) and Markowitz (1976) and show how the performance of myopic mean-variance investors can be improved by constraining them to not exceed growth optimal risk taking.

Markowitz (1976) argues that a *“Kelly-Latané point should be considered the upper limit for conservative choice among E-V efficient portfolios”*². At that time portfolio theory was known, but not yet uncontested. The single-period framework on which much of Markowitz' work is based prompted early criticism on the applicability of portfolio theory to real world situations. Authors such as Hakansson (1971b,c), Bicksler and Thorp (1973) and Rubinstein (1976) argue that given the shortcomings of portfolio theory, rather the growth optimum model and generalized logarithmic utility should be the premier model of financial markets. The fact that portfolio theory was essentially a single period concept and not straightforward to

¹This chapter is based on the homonymous working paper Winselmann (2018).

²Markowitz (1976) uses E for mean return and V for variance of returns

generalize to multi-period problems was already acknowledged by Markowitz (1959), whose seminal publication on portfolio theory includes a much-ignored chapter on the long run. Despite the early criticism, mean-variance analysis has become the cornerstone of portfolio selection. The growth optimum model and generalized logarithmic utility have since been covered with the dust of time. With them, the work on growth optimal strategies of Kelly (1956) and Latané (1959) have disappeared from academic discussion.

Samuelson (1969) analyses discrete time formulations of portfolio choice problems and Merton (1969, 1971) analyses the respective continuous time equivalents. While their works constitute important advancements of lifetime portfolio selection problems, both require restrictive assumptions on utility functions and return distributions. Samuelson assumes isoelastic utility and independence of returns over time and Merton in his continuous time approach obtains explicit solutions only for utility functions with hyperbolic absolute risk aversion and geometric Brownian motion asset prices, which again entails independence of returns over time. Solutions for discrete time multi-period problems when utility functions are not of the isoelastic type or returns are not independent over time have not been obtained. Steinbach (2001) provides an encompassing survey of mean-variance optimization both in single period and multiperiod settings with a distinct focus on how objective and constraints play together in the various settings.

The assumption of independence of returns over time faces severe challenges in real-world applications. Financial time series data for example frequently exhibits conditional heteroscedasticity. A common justification to still use the myopic assumption is that the loss in expected utility from errors caused by additional model complexity and parameter estimation outweigh the gain from theoretically optimal investing (Brandt, 2009; DeMiguel et al., 2009). Theoretically, hedging demands are only zero when investment opportunities are either constant over time or unheadgable, where unheadgable means that the variation in returns is entirely independent of past returns (Brandt, 2009), or when utility is of the logarithmic type (Hakansson, 1971b).

Given the complexity of mean-variance analysis in multi-period settings, it took close to half a century until Li and Ng (2000) presented an analytically optimal solution to multi-period mean-variance problems and an expression of the efficient frontier for multi-period mean-variance problems. Li and Ng achieve their breakthrough result by using embedding techniques to overcome the nonseparability in the sense of dynamic programming of the variance of wealth expression in multi-period mean-variance problems. Further advancements have followed the work of

Li and Ng. Zhu et al. (2004) introduce risk control over bankruptcy and present an analytically optimal solution to a multi-period mean-variance problem in discrete time. Zhu et al. follow the embedding scheme of Li and Ng to solve the variance control problem. Wu and Li (2011) investigate a multi-period mean-variance portfolio selection problem with regime switching and uncertain exit time and present a closed form solution using embedding techniques. Wu and Li (2012) derive the closed form optimal strategy and efficient frontier in the presence of an additional stochastic cash flow component and regime switching. Wu et al. (2014) generalize Wu and Li (2011) to a setting where exit time is dependent on the state of the market. Cui et al. (2014) extend Li and Ng (2000) to a setting with a no shorting constraint which in a discrete time setting implies a piecewise quadratic value function. These and many other related advancements have—at least from a theoretical perspective—substantially advanced our ability to solve multi-period mean-variance problems analytically and optimally.

An equilibrium based approach to solve multi-period mean-variance problems is presented in Basak and Chabakauri (2010) who investigate a dynamic mean-variance problem in a Markovian setting and present an explicit time-consistent solution using dynamic programming and the Hamilton-Jacobi-Bellman equation. Björk et al. (2014) extend the work of Basak and Chabakauri (2010) from constant risk aversion to state dependent risk aversion.

In practice, multi-period models are rarely used for a variety of reasons. In order to determine an optimal strategy for multi-period portfolio selection, return and risk estimates for all periods are required at $t = 0$. Apart from the impracticability and often unavailability of such estimates, multi-period models are computationally intensive. Moreover, such models are frequently not compatible with real world constraints (Kolm et al., 2014). As a consequence, multi-period problems are in academia and industry often modeled as series of myopic single-period decisions, despite the resulting investment decisions not being theoretically optimal (Basak and Chabakauri, 2010).

Markowitz (1976) has pointed to the relevance of the work of Kelly (1956) and Latané (1959) when applying myopic single-period mean-variance analysis to multi-period problems. What he refers to as the “*Kelly-Latané point*” on the efficient frontier is the growth optimal portfolio. The choice of portfolios with, from a single period perspective, higher arithmetic mean returns results in greater volatility of investor portfolio returns over time and less long run investor wealth growth.

This paper contributes to the literature by investigating how the “*Kelly-Latané point*” can be applied to myopic mean-variance investor risk taking decisions in

multi-period asset allocation problems in order to protect long run growth. We include a risk free asset and thus move from the efficient frontier of Markowitz (1952) to the capital market line of Tobin (1958). We present a two asset case with a risk free asset and a risky asset, for which we assume normally distributed returns. We compare the investment decisions of a growth optimal investor with those of a mean-variance optimizer. For both investors the decision variable is the fraction of wealth to invest into the risky asset given the mean and the variance of next period's returns. For the implementation of the growth optimal investor we follow Rotando and Thorp (1992), who have presented an application of the Kelly criterion to the stock market. The Kelly criterion was originally formulated for gambles with binomial outcomes. The generalization of Rotando and Thorp applies the Kelly criterion to the stock market where wealth and stock prices much rather have continuous distributions of outcomes. We find that by constraining the risk taking choice of the myopic mean-variance investor to not exceed the growth optimal level of risk taking, both Sharpe ratio and long run wealth growth of the mean-variance investor can benefit.

3.2 Asset allocation decisions

In our capital growth problem in the presence of uncertainty, investors decide at each decision point how to allocate capital between a risky and a risk free asset. This is equivalent to making investment decisions along the capital market line (Tobin, 1958). We assume risky asset returns at time t to be normally distributed with mean μ_t and standard deviation σ_t . We further assume that mean and standard deviation are the primary return characteristics of interest that investors employ to choose between portfolios. These assumptions fit the so-called "CAPM world" of Markowitz (Markowitz, 1959). In the multi-period problem which we present as an application, the distribution parameters μ_t and σ_t at some future point in time may be different from today³. We do not require the assumption of independence of returns over time since our mean-variance investor by definition is myopic and the growth optimal investor with log utility is indifferent to that assumption.

We model a growth optimal investor and a mean-variance investor to analyze their capital allocation decisions and contrast their performance in a multi-period investment problem. The growth optimal investor considers the entire distribution

³A simple example explains why this fits reality: assume you ask capital market experts in the middle of a tranquil and stable bull market phase for an estimate of next month's expected return and volatility of the S&P500 index. Now, assume a turbulent market phase—say, Lehman Brothers went bankrupt the day before—and you ask the same experts again. You will most likely get a significantly different consensus estimate for both, expected return and volatility.

of returns in his optimization problem at each point in time. Mean-variance analysis as pioneered by Markowitz in contrast has a normal distribution of returns built-in as an assumption. In the utility maximization problem of the mean-variance investor we only find the mean and the variance, but not the probability density function itself. We thus have both investors on a level playing field when we assume normally distributed returns with mean μ_t and standard deviation σ_t as the assumptions and parameters underlying both investors' asset allocation decisions are identical.

In our analysis we follow Rotando and Thorp (1992) and assume discrete returns to be normally distributed. They acknowledge that the log-normal distribution—and thus normally distributed continuous time returns—may provide a superior fit to real-world stock returns, but given that the assumption of log-normal returns makes the computations for the growth optimal investor much more cumbersome, they proceed with normally distributed returns which they find to also be a reasonably close approximation to stock market returns. To ensure that returns are on a meaningful interval, they cap the tails of the normal distribution.

On a more theoretical level, the argument for the log-normal distribution generally is that it allows to aggregate returns over time easily. In contrast, the argument for the normal distribution is that it allows to aggregate normally distributed asset returns to a portfolio return which then itself is normally distributed. This is not the case for log-normal asset returns as the weighted sum of correlated log-normal random variables is not log-normally distributed, but its distribution can only be approximated (Campbell and Viceira, 2002; Lo, 2012). This provides good reason to not work with log-normal returns in mean-variance analysis unless one uses continuous time—which we do not, since continuous time does not fit the constraints of a real-life investor who re-balances his portfolio in discrete time intervals.

In our two asset case the decision variable at time t is the weight in the risky asset $\omega_t \geq 0$, which determines the asset allocation between a risk free asset with certain return $r_{f,t}$ and a risky asset with normally distributed returns with mean μ_t and standard deviation σ_t . A weight of $\omega_t = 0$ represents 100% of investor wealth invested in the risk free asset at time t , and $\omega_t = 1$ represents 100% of investor wealth invested in the risky asset. We assume that leverage in the form of $\omega_t > 1$ is possible, and that the investors can borrow and lend at the rate of return of the risk free asset. We assume the risky asset to have a positive expected excess return over the risk free asset, and thus a short position in the risky asset with $\omega_t < 0$ will never be an optimal decision for any investor. To ease computations, we assume the risk free rate $r_{f,t}$ to be zero. From $r_{f,t} = 0$ follows that the return r_t is equal to the

excess return over the risk free rate and the mean return μ_t is equal to the mean excess return. Further, we assume that the investors incur no trading cost.

3.2.1 Growth optimal investor

The strategy pursued by the growth optimal investor is to maximize the expected log of wealth (Kelly, 1956; Latané, 1959; Samuelson, 1969). By the properties of the log function, the multi-period problem decomposes into a sequence of independent single period problems and the optimal policy is to maximize the expected logarithm of single period wealth growth at each decision point (Mossin, 1968). Any essentially different strategy can analytically be proven to yield inferior long run wealth growth (Breiman, 1961; Hakansson, 1971b).

Let W_t be the investor's wealth at the end of period t . T denotes the time horizon. EU denotes an expected utility function and E is the expectation operator. The investor portfolio discrete return in period t is given by $r_{p,t}$. Equation 3.1 shows, that it is the same for a growth optimal investor to either maximize the expected log of his final wealth W_T , or to maximize the expected log growth rate of his wealth in each period.

$$\begin{aligned}
 EU &= E[\log(W_T)] \\
 &= E\left[\log\left(W_0 * \prod_{t=1}^T (1 + r_{p,t})\right)\right] \\
 &= \log(W_0) + \sum_{t=1}^T E[\log(1 + r_{p,t})]
 \end{aligned} \tag{3.1}$$

We see that the optimal decision policy is myopic. Neither the level of wealth nor the growth rate of wealth of any previous or subsequent period influences the portfolio return $r_{p,t}$ which is solely driven by investor risk taking at time t . Optimal solutions for growth optimal strategies are obtained without recursion and without the need for parameter estimates for more than one future period.

Let $G(\omega_t)$ denote the expected log growth rate of investor wealth in period t as a function of the weight in the risky asset ω_t at time t . With $G(\omega_t) = E[\log(1 + r_{p,t})]$ we obtain

$$\begin{aligned} \max EU &= \max \left[\log(W_0) + \sum_{t=1}^T G(\omega_t) \right] \\ &= \log(W_0) + \max_{\omega_t} \sum_{t=1}^T G(\omega_t) \end{aligned} \quad (3.2)$$

as the initial endowment W_0 is a constant and does not influence the optimal decision policy. The optimal weight in the risky asset of the growth optimal investor can thus be obtained from

$$\operatorname{argmax}_{\omega_t} \left[\log(W_0) + \sum_{t=1}^T G(\omega_t) \right] = \operatorname{argmax}_{\omega_t} \left[\sum_{t=1}^T G(\omega_t) \right]$$

To model our growth optimal investor, we follow Rotando and Thorp (1992), who have outlined an application of the Kelly criterion (Kelly, 1956) to the stock market. The original Kelly criterion is used to achieve maximum growth in an infinite series of risky investments with binomial outcomes. It is asymptotically optimal, meaning that only if the length of a sequence of risky investments converges to infinity it will dominate any essentially different strategy with certainty (Breiman, 1961). Starting from the binomial case of the original Kelly criterion, Rotando and Thorp (1992) present a generalization to a normal distribution of outcomes. In the remainder of this section we briefly summarize those results of their work, which are relevant here and introduce our notation. We will not reproduce the entire derivation and refer the interested reader to the original paper for a more thorough understanding.

From our assumption $r_{f,t} = 0$ follows that the portfolio return in each period is given by $r_{p,t} = \omega_t r_t$. Rotando and Thorp (1992) then find the expected log growth rate of wealth $G(\omega_t)$ to be

$$G(\omega_t) = \int_A^B \log(1 + \omega_t r_t) \psi(r_t) dr_t \quad (3.3)$$

where r_t is assumed to follow a normal distribution density function $\psi(r_t)$ on the interval $[A, B]$, with $A = \mu_t - \phi\sigma_t$ as lower limit and $B = \mu_t + \phi\sigma_t$ as upper limit. ϕ thus determines the width of the interval in which realized returns lie. Rotando

and Thorp use $\phi = 3$ and work with yearly excess return data and rebalancing. The tails of the normal distribution have to be cut in order to obtain meaningful results, given that a positive probability of infinite losses is not only unrealistic for stock market investments, but also incompatible with a growth optimal strategy. For a positive probability of infinite losses, the growth optimal wager will always converge to zero (Rotando and Thorp, 1992). The probability density function of excess returns $\psi(r_t)$ is defined as

$$\psi(r_t) = \begin{cases} h_t + \frac{1}{\sqrt{2\pi\alpha_t^2}} e^{-(r_t-\mu_t)^2/2\alpha_t^2} & \forall A \leq r_t \leq B \\ 0 & \forall r_t < A \vee r_t > B \end{cases} \quad (3.4)$$

where h_t and α_t are required to preserve the desired properties of the probability density function $\psi(r_t)$ in the presence of cut tails. The parameter h_t is needed to keep the value of the integral under the normal distribution density function $\int_A^B \psi(r_t) dr_t$ equal to one. We use h_t to add back on the interval $[A, B]$ the area that was cut in the tails. Adding h_t is comparable to adding the cut tail area back in the form of a rectangular box under the normal distribution density function on the interval $[A, B]$. The cut tail area is given by $1 - \int_A^B \mathcal{N}(\mu_t, \alpha_t)$. Excess returns are defined to be normally distributed with mean μ_t and standard deviation σ_t . Cutting the tails reduces the observed standard deviation. In order to preserve the desired standard deviation σ_t , Rotando and Thorp introduce α_t with $\alpha_t > \sigma_t$. They numerically determine the value of α_t for which excess returns r_t on the interval $[A, B]$ have the observed standard deviation σ_t . This is the case when

$$\sigma_t^2 = \int_A^B r_t^2 \psi(r_t) dr_t - \mu_t^2$$

as the expected value of squared excess returns $E[r_t^2]$ can be written as an integral with respect to the underlying probability measure and $\psi(r_t)$ is the probability density function. Having defined α_t , we can compute h_t from

$$h_t * (B - A) = 1 - \int_A^B \mathcal{N}(\mu_t, \alpha_t)$$

where $(B - A)$ can be imagined as the width of a rectangle and h_t as the height. Applying the probability density function $\psi(r_t)$ from equation 3.4 to equation 3.3 yields the formulation for the expected log growth rate of wealth $G(\omega_t)$ in equation

3.5, which is the objective function that a growth optimal investor seeks to maximize in each period.

$$G(\omega_t) = \int_A^B \log(1 + \omega_t r_t) \left[h_t + \frac{1}{\sqrt{2\pi\alpha_t^2}} e^{-(r_t - \mu_t)^2 / 2\alpha_t^2} \right] dr_t \quad (3.5)$$

The integration required to set $G'(\omega_t) = 0$ and to determine the growth optimal weight in the risky asset is non-elementary and cannot be done explicitly (Rotando and Thorp, 1992). We numerically optimize $G(\omega_t)$ with respect to ω_t and obtain the value ω_t for which $G(\omega_t)$ has its maximum. We call this optimum the growth optimal weight in the risky asset ω_t^{go} , where *go* is used as a superscript to denote growth optimal, in contrast to *mv* which we use for the mean-variance investor. At each decision point, which is the beginning of each period t , we determine the asset allocation of the growth optimal investor, expressed as ω_t^{go} , given some μ_t and σ_t .

3.2.2 Mean-variance investor

The mean-variance investor derives utility from gains in wealth, but suffers disutility from losses and volatility. Preferences for mean-variance investors are either assumed to be given as a function of mean and variance or they are analytically derived from expected utility theory. Both approaches are taken in the literature and good overviews are available elsewhere (Campbell and Viceira, 2002; Brandt, 2009; Markowitz, 2014).

Let $r_{p,t}$ be the return on the investor portfolio in period t and $\sigma_{p,t}$ be the standard deviation of investor portfolio returns. Let γ denote the level of risk aversion of the investor. We assume the mean-variance preferences of the investor to be described by the following expected utility function.

$$EU_t = E[r_{p,t}] - \frac{\gamma}{2} \sigma_{p,t}^2 \quad \forall \gamma > 0 \quad (3.6)$$

The investor thus acts myopic and maximizes his expected utility EU_t in each period t according to equation 3.6, which is a common textbook function for expected utility. The mean-variance investor's expected utility maximizing weight in the risky asset is known to be given by (Campbell and Viceira, 2002; Alexander, 2008)

$$\omega_t^{mv} = \frac{\mu_t - r_{f,t}}{\gamma \sigma_t^2} \quad (3.7)$$

where μ_t is the expected return on the risky asset and σ_t is the standard deviation. We use ω_t^{mv} with mv as a superscript to denote the optimal value for the mean-variance investor. From equation 3.7 we see, that the mean-variance investor's weight in the risky asset depends on the expected excess return over the risk free rate, and not on the total return of the assets. The nominator $\mu_t - r_{f,t}$ represents the expected excess return. Throughout this paper we assume a risk free rate of zero, $r_{f,t} = 0$, and thus the expected return μ_t is equivalent to the expected excess return.

The expected utility function in equation 3.6 and the resulting optimal weight in the risky asset entail quadratic von-Neumann-Morgenstern utility, which may be considered undesirable as it implies increasing absolute risk aversion (Brandt, 2009). Notwithstanding, it is convenient and commonly used in academia and practice. Either, it is assumed to describe the preferences of an investor exactly⁴, or some different investor preferences are stated, but a utility function similar to equation 3.6 is used as an approximation to investor preferences. Markowitz (1959) in his seminal work on portfolio theory has already suggested such approximations and concluded, that quadratic approximations work well for returns in the range of -30% and $+40\%$. A substantial body of research has evolved on mean-variance approximations to expected utility. Levy and Markowitz (1979) is one of the most prominent works in support of it. They find that expected utility approximations based on quadratic utility work very well especially for power utility, which is also confirmed by Pulley (1983). Kallberg and Ziemba (1983) generalize this finding for a broader set of utility functions if time horizons are short. Markowitz (2014) reinforces his earlier findings and their validity. While he embraces the advantages of other utility functions to describe investor preferences, he still emphasizes that these other utility functions can be well approximated with quadratic utility, especially as the translations of other utility functions into mean-variance formulations often come with restrictive assumptions on the return distribution or require their own kind of approximations. Consequently, mean-variance preferences of the form stated in equation 3.6 are a common approach in practice and academia to obtain mean-variance approximations to expected utility. Our model of the mean-variance investor thus is consistent with formulations commonly used in practice and academia.

⁴See for example Markowitz (1959); Brandt and Santa-Clara (2006); Sun et al. (2006); DeMiguel et al. (2009); Basak and Chabakauri (2010) as well as most standard textbooks.

3.3 Theoretical considerations

Our aim is to employ the growth optimal weight in the risky asset as an upper limit for mean-variance investor risk taking to protect long run wealth growth and performance. First, we discuss why higher than growth optimal risk taking can be harmful. Second, in order to understand the asset allocation decisions of both investors, we present their respective investment decisions as functions of the mean excess return with fixed standard deviation and as functions of the standard deviation with fixed mean excess return. Third, we present the investment decisions of both investors in mean-variance space and highlight where excessive risk taking occurs.

3.3.1 Overbetting

A growth optimal investor acts to maximize expected long run wealth growth. Breiman (1961) proves that the final wealth W_T^* accumulated under a growth optimal strategy Λ^* will asymptotically be infinitely larger than the final wealth W_T accumulated under any essentially different strategy Λ , which is not asymptotically close to Λ^* . Let a non-terminating growth optimal strategy Λ^* lead to wealth W_T^* after T plays and let an essentially different strategy Λ lead to wealth W_T . Then almost surely there exists a limit for W_T^*/W_T and $\lim_{T \rightarrow \infty} (W_T^*/W_T) = \infty$ (Breiman, 1961).

The growth optimal weight in the risky asset ω_t^{go} is always the weight that leads to the asymptotically highest long run wealth growth. Any lower ω_t choice will lead to a wealth path with a lower variance but also lower final wealth. Any higher ω_t choice, which literature frequently refers to as overbetting, will lead to a wealth path with a higher variance but lower long run wealth growth (MacLean et al., 1992). A special case of higher ω_t choice constitutes $2\omega_t^{go}$. It can be proven analytically that using twice the growth optimal weight in the risky asset leads to high fluctuations of wealth around the initial wealth with an expected long run growth rate of zero (MacLean et al., 2010c).

The detrimental effect of overbetting in a multi-period context has already been mentioned at the outset of portfolio theory. Markowitz (1959) in his seminal monograph on portfolio theory discusses the difference between arithmetic and geometric returns and explains how a series of investments with a positive expected arithmetic return can lead to virtually certain bankruptcy in the long run.

To illustrate the argument of Markowitz (1959), let us assume there are two investment opportunities X and Y, each with two possible outcomes of equal probability. Investment X either has a return of -25% or +75%, giving it an expected

return of +25%, and investment Y either has a return of -50% or +150%, giving it an expected return of +50%. Investment X has a geometric mean return of $\sqrt{(1 - 25\%)(1 + 75\%)} - 1 = 14.6\%$ and investment Y has a geometric mean return of $\sqrt{(1 - 50\%)(1 + 150\%)} - 1 = 11.8\%$. The lower geometric mean return of the investment with higher arithmetic mean return highlights the damage that variance in arithmetic returns does to long term wealth accumulation, which is correctly measured by geometric returns (Ziembra, 2005). This effect is more obvious if one includes an investment that may suffer a total loss. Assume there also is investment Z, which has either a return of -100% or +200%. Albeit the arithmetic mean is +50%, the geometric mean is $\sqrt{(1 - 100\%)(1 + 200\%)} - 1 = -100\%$ and if all money is re-investment indefinitely, bankruptcy is virtually certain. A single occurrence of the -100% outcome wipes out all capital.

From the examples above, we see that variance in arithmetic discrete returns reduces wealth growth in multi-period asset allocation problems. In the heat of the academic discussion about the advantages and disadvantages of using the Kelly criterion and growth optimal strategies (Samuelson, 1969; Hakansson, 1971b; Samuelson, 1971; Roll, 1973; Rubinstein, 1976), Markowitz has therefore asserted that it is reasonable to warn *“an investor against choosing E-V efficient portfolios with higher E and V but smaller $E \log(1 + r)$ [...] on the grounds that such higher E-V combinations have greater variability in the short run and less return in the long run”*⁵ (Markowitz, 1976). In the same publication Markowitz refers to the point with approximate maximum $E \log(1 + r)$ as the *“Kelly-Latané point”*, which *“should be considered the upper limit for conservative choice among E-V efficient portfolios”*.

The Kelly criterion and growth optimal strategies might provide valuable information to a myopic mean-variance optimizer. Namely, that if he chooses portfolios that are riskier than the growth optimal portfolio, he might be maximizing his expected utility in a myopic one period context, but at the same time he might accept variance in portfolio returns over time that reduce the long term growth he realizes. This would be straightforward and not of interest, if single period and multi-period problems were clearly separated and applied in a mathematically correct way in their respective domains. Recalling the frequent modeling of multi-period mean-variance problems as series of myopic single-period decisions (Basak and Chabakauri, 2010), the remark of Markowitz (1976) becomes relevant. Markowitz’ argument was based on the efficient frontier, but it is just as valid for the capital market line that we obtain when a risk free asset is introduced.

⁵Markowitz uses E for expected excess return and V for variance of returns.

3.3.2 The weight in the risky asset

In this section, we analyze for which risky asset returns a myopic mean-variance investor chooses an asset allocation that is beyond growth optimal. We first discuss the asset allocation decisions of the growth optimal investor and the mean-variance investor individually, and then compare both investors. This makes us understand, what combinations of next period risky asset mean return μ_t and variance σ_t^2 lead the mean-variance investor to take a higher weight in the risky asset than growth optimal. We show that myopic mean-variance optimizers particularly take excessive risk when the expectation for next period risky asset returns is favorable, i.e. when the expected excess return on the risky asset μ_t is high and variance of risky asset excess returns σ_t^2 is low. As before, we assume a risk free rate of zero. We use values for risky asset mean excess return on the interval [1%, 15%] and values for standard deviation of risky asset excess returns on the interval [5%, 20%]. Historic yearly stock market excess returns and standard deviation commonly are in this range. The yearly excess return on U.S. stocks over the last 30 years has been on average 8.4% and the standard deviation has been 15%⁶.

3.3.2.1 Growth optimal investor

We have outlined the modeling of the growth optimal investor in section 3.2.1. Underlying assumption of our model is that discrete returns are normally distributed on the interval $[A, B]$, with $A = \mu_t - \phi\sigma_t$ and $B = \mu_t + \phi\sigma_t$. The required parameters to fully specify the distribution of next period risky asset returns and to determine the growth optimal weight in the risky asset are the expected mean excess return μ_t , the standard deviation σ_t , and ϕ , which determines the width of the permissible interval for realized risky asset returns. We define $\phi = 4$ for our analysis, meaning that a deviation of $4\sigma_t$ to the left and right of μ_t is the minimum and maximum permissible risky asset excess return realization.

To illustrate the risk taking behavior of the growth optimal investor, we set the mean excess return on the risky asset to 8.4% and calculate the growth optimal weight in the risky asset for 50 different values of standard deviation of risky asset returns on the interval [5%, 20%]. The resulting weights in the risky asset are shown in figure 3.1a. Weights in the risky asset for $\phi = \{3, 4, 5\}$ are shown. The axis showing the weight in the risky asset is log scaled to better visualize results. The growth optimal risky asset weight in figure 3.1a decreases as the value of ϕ increases, i.e. as the tails of the distribution get longer. This is due to the possibility of higher

⁶U.S. research returns data (French, 2016).

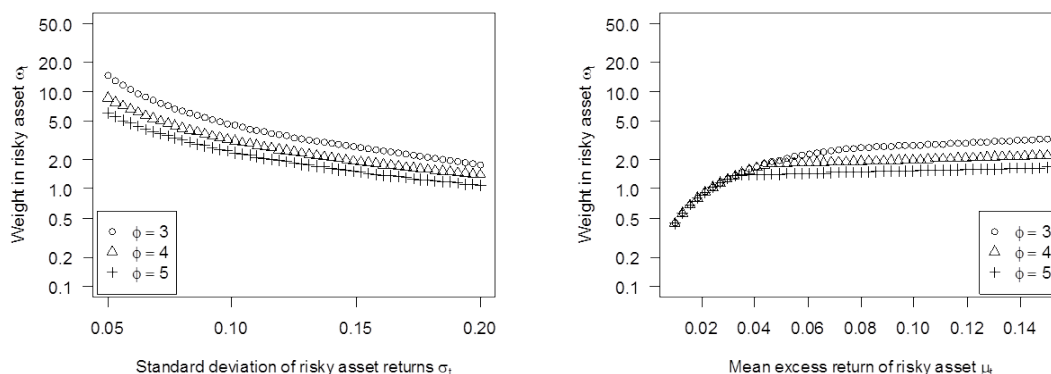
losses, which have a more detrimental effect on wealth growth than the positive effect of the possibility of higher gains. Also, the maximum weight in the risky asset ω_t for which the possibility of bankruptcy yet is zero decreases as ϕ increases. In figure 3.1a we see that ceteris paribus as the standard deviation of risky asset returns increases, the growth optimal risk taking decreases and asymptotically approaches zero. The growth optimal risk taking asymptotically increases towards infinity as the standard deviation of risky asset returns decreases. It reaches infinity when the interval $[A, B]$ is strictly non-negative, i.e. $\mu_t \geq \phi\sigma_t$ and thus for all return realizations possible on $\psi(r_t)$ we have $r_t \geq 0$. This is intuitive as it is certainly optimal to employ infinite leverage when the worst case excess return is zero.

Figure 3.1b shows the growth optimal weight in the risky asset for 50 different values of risky asset mean excess return on the interval $[1\%, 15\%]$ and 15% standard deviation. Again we see the optimal weight in the risky asset decrease as ϕ increases, making the interval wider on which realized returns can lie. We see the optimal weight in the risky asset increase as risky asset excess return on the interval $[1\%, 15\%]$ increases. The weight in the risky asset asymptotically increases towards infinity as the risky asset mean excess return increases towards the level where no negative risky asset excess return realizations have a positive probability, i.e. where $\mu_t \geq \phi\sigma_t$ and thus $r_t \geq 0$. The growth optimal weight in the risky asset decreases to zero as the mean excess return on the risky asset decreases to zero.


3.3.2.2 Mean-variance investor

Relevant input for the mean-variance investor are three parameters: the risky asset mean excess return μ_t , the standard deviation of returns on the risky asset σ_t , and the risk aversion parameter γ . Based on these three parameters, the mean-variance investor described in section 3.2.2 determines his utility maximizing allocation of assets between the risk free asset and the risky asset with normally distributed returns.

To illustrate the risk taking behavior of the mean-variance investor, we again set the mean excess return on the risky asset to 8.4% and calculate his weight in the risky asset for 50 different values of standard deviation of risky asset returns on the interval $[5\%, 20\%]$. The resulting weights in the risky asset are shown in figure 3.1c. Weights in the risky asset for multiple values of risk aversion γ are shown. We see the weight in the risky asset decrease as risk aversion increases. As standard deviation of risky asset returns on the interval $[5\%, 20\%]$ increases, we observe the mean-variance investor risk taking decrease. Equation 3.7 in section 3.2.2 shows that



(a) Weight in risky asset chosen by growth optimal investor for risky asset mean excess return of 8.4% and 50 values of risky asset std. dev. on the interval [5%, 20%] (b) Weight in risky asset chosen by growth optimal investor for risky asset std. dev. of 15% and 50 values of risky asset mean excess return on the interval [1%, 15%]



(c) Weight in risky asset chosen by mean-variance investor for risky asset mean excess return of 8.4% and 50 values of risky asset std. dev. on the interval [5%, 20%] (d) Weight in risky asset chosen by mean-variance investor for risky asset std. dev. of 15% and 50 values of risky asset mean excess return on the interval [1%, 15%]

Figure 3.1: Investor’s chosen weight in the risky asset for various combinations of risky asset mean excess return and standard deviation; growth optimal investors for different values of return distribution width ϕ and mean-variance investors with different values of risk aversion γ are shown

the mean-variance investor’s weight in the risky asset is, ceteris paribus, inversely proportional to the variance of risky asset returns σ_t^2 .

Figure 3.1d shows the weight in the risky asset of the mean-variance investor for 50 different values of risky asset mean excess return on the interval [1%, 15%] and 15% standard deviation of risky asset returns. Again we see the weight in the risky asset decrease as the risk aversion parameter γ increases. We see risk taking increase as risky asset excess return on the interval [1%, 15%] increases. Equation 3.7 shows that ceteris paribus the mean-variance investor’s weight in the risky asset is proportional




Figure 3.2: Surface plot of growth optimal investor (black grid) and mean-variance investor (grey surface, $\gamma = 3$) weight in the risky asset for different values of risky asset mean excess return μ_t and standard deviation σ_t

to the mean excess return. As the mean excess return increases towards infinity *ceteris paribus*, the weight in the risky asset also increases to infinity. The weight in the risky asset decreases to zero as risky asset mean excess returns decrease to zero.

3.3.2.3 Comparison of risk taking preferences

We now compare the risk taking of both investors presented above. Our aim is to better explain, for which combinations of risky asset mean excess return, standard deviation and mean-variance investor risk aversion the mean-variance investor exhibits higher than growth optimal risk taking. We do not simply plot the risk taking depicted in figure 3.1 into combined graphs. While that would indeed show us some data points where the mean-variance investor exceeds growth optimal risk taking, we would yet gain little understanding which combinations of risky asset mean excess return and standard deviation generally lead to higher than growth optimal risk taking. As the required numerical work to obtain the growth optimal weight in the risky asset precludes the formulation of a closed form solution, we resort to a surface plot that visualizes investor risk taking.

We present a three dimensional graph with mean excess returns and standard deviation in the x-y plane and investor risk taking as the z-axis. This graph allows us to visualize for what combinations of mean excess return and standard deviation the mean-variance investor takes more risk than the growth optimal. We plot two surfaces in the graph, where one surface shows the mean-variance investor's choice of

weight in the risky asset and the other surface shows the growth optimal investor's choice of weight in the risky asset, both as a function of the mean excess return and the standard deviation of returns. To obtain the required data, we compute the weights in the risky asset of both investors for 50 different values of excess returns on the interval $[1\%, 15\%]$ and 50 different values of standard deviations on the interval $[5\%, 20\%]$, which is 2,500 data points for each investor. Figure 3.2 presents both investors' risk taking choices. The mean-variance investor's weight in the risky asset is plotted as a grey surface and the growth optimal investor's weight in the risky asset is plotted as a black grid. We see that an area exists, where the mean-variance investor's weight in the risky asset exceeds that of the growth optimal investor. This area of mean-variance investor overbetting covers favorable market conditions with high mean excess return and low standard deviation.

To simplify the result, let us abstract from the exact weight in the risky asset that the investors take. Instead, let us only concentrate on whether or not the mean-variance investor exhibits higher than growth optimal risk taking. The z-axis, which shows the weight in the risky asset in figure 3.2, therefore collapses into a binary result: overbetting or no overbetting. In figure 3.3 we show this simplified result for four mean-variance investors. As above, the mean excess return is on the interval $[1\%, 15\%]$ and the standard deviation is on the interval $[5\%, 20\%]$.

We have computed the weight in the risky asset for four mean-variance investors with different degrees of risk aversion. For each mean-variance investor we have shaded in grey the area where he allocates a higher fraction of wealth to the risky asset than the growth optimal investor. We find that the low risk aversion mean-variance investor with $\gamma = 1$ overbets for all depicted combinations of mean excess return and standard deviation. For the mean-variance investors with $\gamma = \{2, 3, 4\}$ we note overbetting in the upper left of the graph, which means for combinations of higher mean excess return and lower standard deviation. In figure 3.3 we see that the area in which overbetting occurs shrinks when the mean-variance investor risk aversion parameter increases. Initially, it moves further to the upper left. For $\gamma = 4$, it finally dwindles to a small band that roughly lies between the point with 10% mean excess return and 5% standard deviation and the point with 15% mean excess return and 7.5% standard deviation. Taking a closer look at the three dimensional plot in figure 3.2 before, we see that at the very left of the graph the extent by which the mean-variance investor overbets is not at its maximum. Rather, for a mean excess return of 15%, it starts to reduce again as standard deviation finally decreases towards 5%. This is driven by the growth optimal investor's weight in the risky asset increasing towards infinity as no negative return realization on $\psi(r_t)$ in

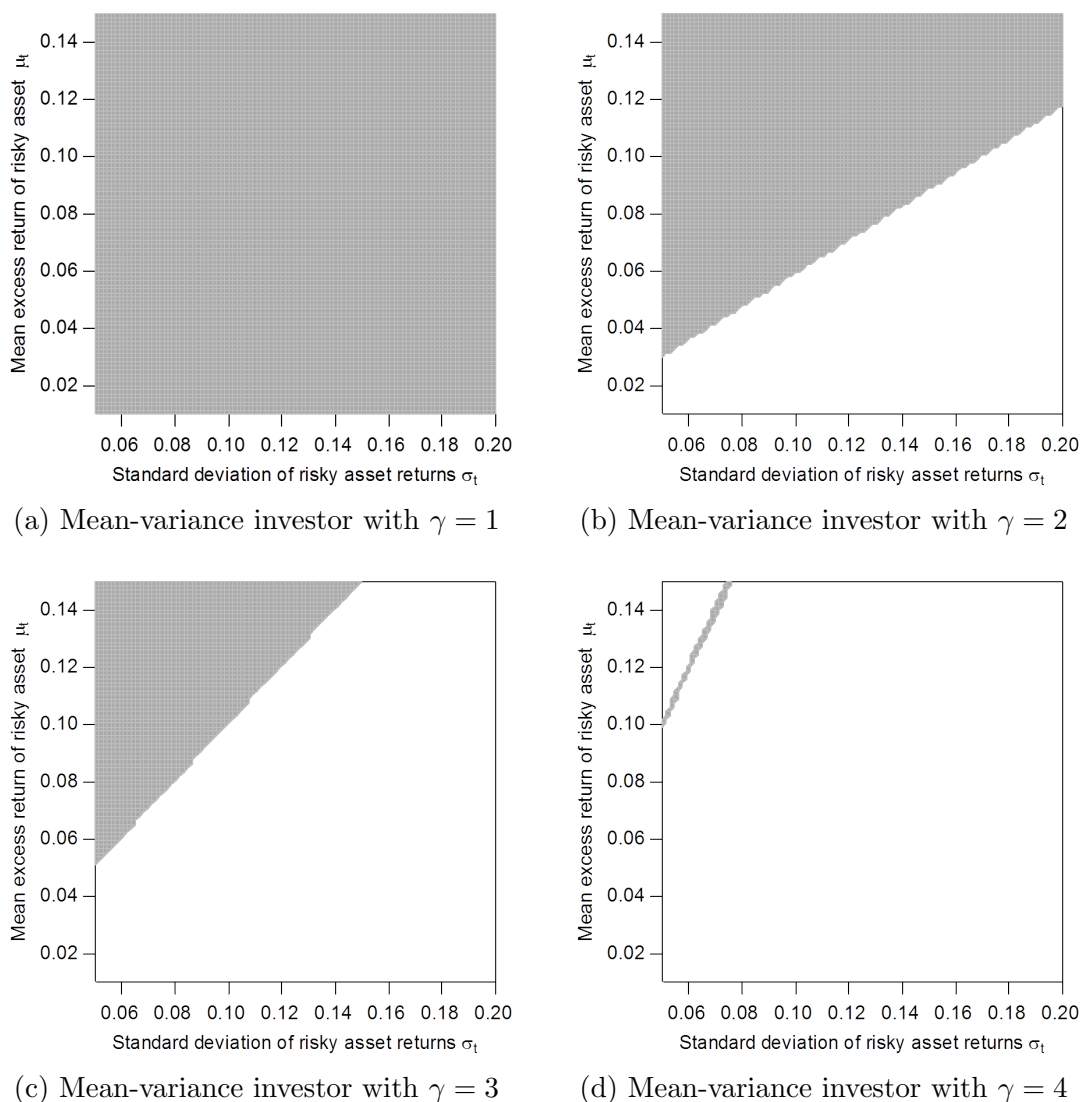


Figure 3.3: Overbetting of mean-variance investors with four different values of risk aversion γ ; grey color indicates areas of risky asset mean excess return and standard deviation where mean-variance investors exhibit higher than growth optimal risky asset weight

the interval $[A, B]$ is possible. This effect takes place in extremely favorable market conditions with very low standard deviation of returns and comparatively high mean returns, which at least for stock markets we deem unrealistic.

Note that for $\gamma = 3$ the grey shaded area in figure 3.3c covers the same area of mean excess return and standard deviation combinations, for which in figure 3.2 the grey surface lies above the black grid. Figure 3.3c and figure 3.2 both depict the overbetting of the $\gamma = 3$ mean-variance investor.

3.4 Application example

In section 3.3.1 we have discussed the theoretical effects of overbetting and in section 3.3.2.3 we have shown that mean-variance investor overbetting indeed occurs if risk aversion is low. In this section we investigate with simulation studies the effect that overbetting has on wealth accumulation and investor performance. We demonstrate in Monte Carlo simulations the negative effect of overbetting on wealth growth and investor performance in multi-period investing.

First, we present a simulation with constant investment opportunities where the investors choose an asset allocation expressed in fractions of wealth that remain constant over all periods. We find that higher than growth optimal risk taking can lead to high final wealth outcomes if time horizons are finite, but also results in frequent bankruptcy of mean-variance investors and in lower Sharpe ratios.

Second, we present a simulation with stochastic investment opportunities which we deem closer to real life multi-period investment problems. We again find that the use of the growth optimal weight in the risky asset as an upper limit for mean-variance investor risk taking does increase the Sharpe ratio and can protect mean-variance investors from bankruptcy.

3.4.1 Constant investment opportunities

We assume a risky asset with a return of 8.4% and a standard deviation of returns of 15.0%. These values correspond to the last thirty years of yearly stock market returns observed by French (2016) in his U.S. research returns data set. We assume the return distribution of the risky asset to be given by $\psi(r_t)$ in equation 3.4 in section 3.2.1. We cut the tails at $\pm 4\sigma$, which we deem a reasonable interval for stock market returns. We run 1,000 Monte Carlo simulation with 100 periods each and let our growth optimal and mean-variance investors allocate assets between the risky asset and the risk free asset. We compute average final wealth and average Sharpe ratios (Sharpe, 1966, 1994) of the investors to compare their performance.

For the growth optimal investor we find the weight in the risky asset to be 1.94, meaning that a growth optimal investor invests 194% of his wealth in the risky asset and borrows an amount equal to 94% of wealth at the risk free rate to finance his position in the risky asset. The mean-variance investor asset allocation depends on his level of risk aversion. His investment into the risky asset is shown in table 3.1. We see that the mean-variance investor with $\gamma = 2$ is closest to the growth optimal level of risk taking at 188% and all investors with $\gamma < 2$ take more risk than growth optimal.

Table 3.1: Weight of the risky asset in mean-variance investor portfolios in percent of wealth; investment levels $> 100\%$ are enabled by borrowing at the risk free rate

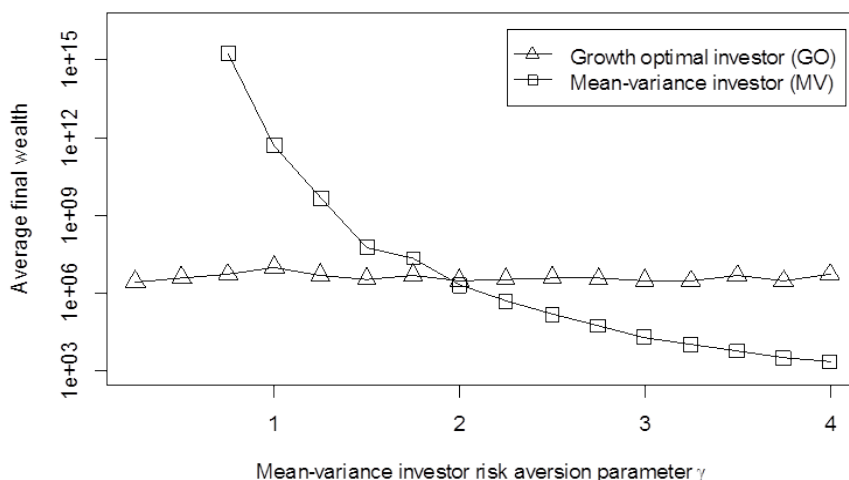
Risk aversion γ	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0
Wealth invested in risky asset	751%	375%	250%	188%	150%	125%	107%	94%

The resulting average final wealth and average Sharpe ratios that our investors achieve during the simulation are shown in figure 3.4. On the x-axis we have the mean-variance investor risk aversion parameter $\gamma = \{0.25, 0.5, \dots, 4.0\}$. For each of these values we have conducted the above described 1,000 Monte Carlo simulations with 100 periods each. On the y-axis we have either average final wealth or average Sharpe ratios.

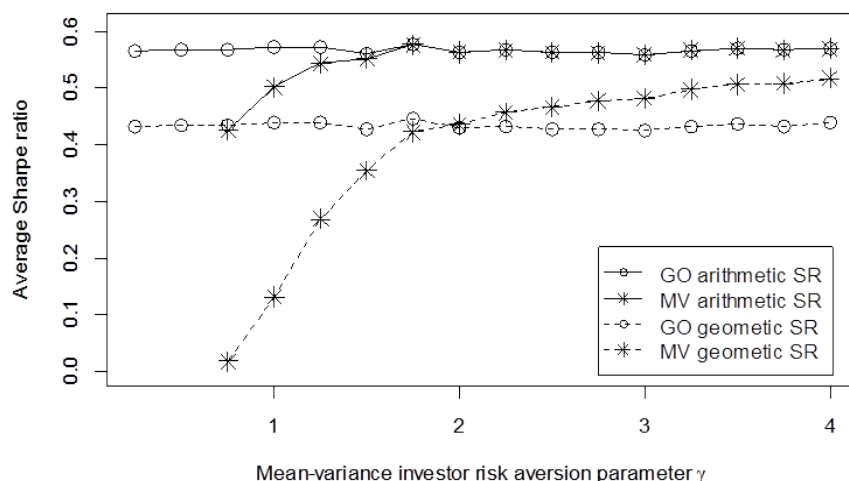
Average final wealth values over 1,000 simulation runs are shown in figure 3.4a. We see the growth optimal investor achieve final wealth levels slightly above one million after 100 periods, starting from an initial wealth of 1 unit. The fluctuations in the growth optimal investor's average final wealth are due to different sets of random return realizations being drawn in the simulation runs. The mean-variance investor's average final wealth is low for high values of risk aversion and increases as the risk aversion decreases. We see that the average final wealth for overbetting mean-variance investors with $0.75 \leq \gamma \leq 1.75$ surpasses that of the growth optimal investor. The aggressively overbetting mean-variance investors with $\gamma \leq 0.5$ have negative levels of average final wealth which in figure 3.4a are omitted.

We know that the growth optimal strategy Λ^* asymptotically dominates any essentially different strategy Λ in terms of wealth accumulation as the time horizon converges to infinity. In contrast, figure 3.4 suggests that some overbetting mean-variance investors surpass the average final wealth of the growth optimal investor. A more detailed investigation of the high average final wealth of the overbetting mean-variance investors reveals that not all investors survive. Table 3.2 shows details. The mean-variance investor with $\gamma = 2$ does not experience bankruptcy during any of the 1,000 simulation runs. All overbetting mean-variance investors of lower risk aversion experience bankruptcy⁷ at least during some simulation runs. The mean variance investors with $\gamma \leq 0.5$ experience bankruptcy during all 1,000 simulation runs. Surviving mean-variance investors with $\gamma < 2$ achieve extremely high maximum final wealth levels. We define maximum final wealth as the highest final wealth that an investor achieved over our 1,000 simulation runs. The very high levels of

⁷We define an investor to be bankrupt if $W_t \leq 0$ for any $t \in \{1, \dots, T\}$ and in such case regard the first negative wealth occurrence as the respective investors final wealth.



(a) Average final wealth of growth optimal and mean-variance investor



(b) Average arithmetic and geometric Sharpe ratio (SR) for growth optimal investor (GO) and mean-variance investor (MV)

Figure 3.4: Average final wealth (a) and average arithmetic and geometric Sharpe ratio (b) of growth optimal and mean-variance investors with different values of risk aversion $0 < \gamma \leq 4$ over 1,000 simulation runs with 100 periods each

final wealth that some of the overbetting mean-variance investors achieve result in average final wealth levels that surpass those of the growth optimal investor.

Table 3.2 shows that the final wealth levels of our investors are extremely divergent. Minimum and maximum final wealth levels over our 1,000 simulations lie far apart. These results are in line with other studies that investigate long term capital growth in multi-period problems (Ziemba and Hausch, 1985; MacLean et al., 1992; MacLean and Ziemba, 1999; MacLean et al., 2010a, 2011). Mean-variance investor

Table 3.2: Final wealth and bankruptcies of mean-variance investors and growth optimal investors in 1,000 simulation runs with both, time horizons of 100 and 200 periods length

	Investor final wealth over 1,000 simulation runs with ...							
	...100 periods each				...200 periods each			
	Min	Avg	Max	Bankrupt	Min	Avg	Max	Bankrupt
MV with $\gamma = 0.5$	NA	NA	NA	1000/1000	NA	NA	NA	1000/1000
MV with $\gamma = 1.0$	NA	5.0e+11	3.5e+14	588/1000	NA	1.5e+19	1.7e+22	869/1000
MV with $\gamma = 1.5$	NA	5.8e+7	5.1e+9	60/1000	NA	3.7e+16	3.1e+19	115/1000
MV with $\gamma = 2.0$	22.6	2.0e+6	2.4e+8	0/1000	6.5e+4	2.7e+12	6.5e+14	0/1000
MV with $\gamma = 2.5$	28.3	1.5e+5	1.5e+7	0/1000	8.8e+3	6.9e+10	5.3e+13	0/1000
MV with $\gamma = 3.0$	15.5	2.0e+4	1.5e+6	0/1000	4.4e+3	4.1e+8	5.5e+10	0/1000
MV with $\gamma = 3.5$	19.0	6.0e+3	4.4e+5	0/1000	3.4e+3	3.9e+7	8.7e+9	0/1000
MV with $\gamma = 4.0$	9.0	2.4e+3	6.8e+4	0/1000	2.2e+3	3.7e+6	1.3e+8	0/1000
Growth optimal	18.1	4.3e+6	1.4e+9	0/1000	6.9e+4	1.2e+13	1.4e+16	0/1000

NA refers to bankruptcies or negative values of wealth; MV stands for mean-variance investor

risk aversion of $\gamma \leq 4$ is low, but may well be considered realistic (Markowitz et al., 1994).

Simulations with longer time horizons show that an increasing fraction of overbetting mean-variance investors experiences bankruptcy. Table 3.2 shows that the mean-variance investors with $\gamma = 1.0$ experiences bankruptcy in 588 of the 1,000 simulation runs when the time horizon is 100 periods. If the time horizon is 200 periods, bankruptcy occurs in 869 of the 1,000 simulations. As the time horizon increases towards infinity, the fraction of overbetting mean-variance investors that experience bankruptcy asymptotically increases towards 100%. From the theory of growth optimal strategies we know, and simulations confirm, that for infinite time horizons no investor who follows an essentially different strategy exceeds the final wealth of the growth optimal investor.

We use Sharpe ratios to assess the risk adjusted performance of our investors. For each simulation run we compute the Sharpe ratio that the investor achieved over the 100 periods. We divide mean excess return of investor portfolio returns over 100 periods by standard deviation. We follow related literature in that we show both, the Sharpe ratio calculated from arithmetic means and arithmetic standard deviations and the Sharpe ratio calculated from geometric means and geometric standard deviations (Ziemba, 2005). Ziemba finds the conclusions of his research, such as the relative ranking of funds by Sharpe ratio, to be the same for both, the arithmetic and the geometric Sharpe ratio. He notes that the difference between arithmetic and geometric Sharpe ratio is a function of return volatility. The geometric mean

of investor portfolio returns $\mu_{p,g}$, where g denotes geometric, over time horizon T is calculated as

$$\mu_{p,g} = \left(\prod_{t=1}^T (1 + r_{p,t}) \right)^{1/T} - 1 = \exp \left(\frac{1}{T} \sum_{t=1}^T \log(1 + r_{p,t}) \right) - 1 \quad (3.8)$$

The geometric standard deviation $\sigma_{p,g}$ over time horizon T is calculated as

$$\sigma_{p,g} = \exp \left(\sqrt{\frac{\sum_{t=1}^T (\log \frac{1+r_{p,t}}{1+\mu_{p,g}})^2}{T}} \right) - 1 \quad (3.9)$$

which entails that the results are not real numbers if losses in excess of 100% occur. We therefore set the geometric Sharpe ratio of bankrupt investors to zero⁸. While Ziemba (2005) argues that the geometric Sharpe ratio is the more appropriate measure of performance, it clearly has the disadvantage of being more difficult to understand. Geometric standard deviation is not an intuitive concept.

Figure 3.4b shows the average arithmetic and average geometric⁹ Sharpe ratios that our investors achieved in 1,000 simulation runs of 100 periods. When arithmetic means and standard deviations are used, we find the Sharpe ratios of the growth optimal investor and the more risk averse mean-variance investors with $\gamma \geq 2$ to be equal. This is explained by both investors choosing constant points on the capital market line, where the risk/return trade-off measured by the Sharpe ratio is linear. This linearity comes to an end when overbetting mean-variance investors experience bankruptcy. Bankrupt investors experience high volatility, but end with low mean returns. Mean-variance investors that experience bankruptcy have low or negative arithmetic mean Sharpe ratios and cause the decline in average Sharpe ratios which we see in figure 3.4b for $\gamma \leq 1.75$ ¹⁰.

For geometric Sharpe ratios, we see in figure 3.4b that they are generally lower, which is consistent with related literature (Ziemba, 2005). The reason is that, first, geometric mean returns are generally smaller than the corresponding arithmetic mean returns. Second, the geometric standard deviation may even be higher than

⁸Most of the bankrupt investors do not end with a wealth of exactly zero but instead end with negative values of wealth. The geometric Sharpe ratios thus should intuitively rather be negative. As geometric means of time series with losses in excess of 100% cannot be expressed in real numbers, we set those values to zero instead as the geometric Sharpe ratio approaches zero as wealth approaches zero. One might argue that reality is even worse.

⁹Each shown result is an arithmetic average over the 1,000 geometric Sharpe ratios obtained from our 1,000 simulation runs.

¹⁰For $\gamma \leq 1.5$ the difference is visible in figure 3.4b with the naked eye, for $\gamma = 1.75$ it is in the third digit after the comma only.

the arithmetic standard deviation since high losses translate into high absolute values of log returns. This explains why the geometric Sharpe ratios of low risk aversion mean-variance investors are particularly low. Also, this explains why the geometric Sharpe ratios of the growth optimal investor are lower than those of more risk averse mean-variance investors with $\gamma \geq 2$.

From this simulation we learn that mean-variance investor overbetting leads to Sharpe ratios lower than those of the growth optimal investor. In addition, we learn that mean-variance investor overbetting in a series of favorable investment opportunities can lead to near-certain ruin. Nevertheless, we note that for finite time horizons some overbetting mean-variance investors—in our case those with $0.75 \leq \gamma \leq 1.75$ —can achieve higher average final wealth values than the growth optimal investor.

3.4.2 Stochastic investment opportunities

Investment opportunities in real life multi-period investment situations over long time horizons are not constant. For example, stock returns during the 15 years from 2000-2014 had an annual mean excess return of 4.5% and a standard deviation of 17.8%, while stock returns during the 15 years from 1984-1999 had an annual mean excess return of 12.4% and a standard deviation of 10.7%¹¹. Stochastic investment opportunities that change randomly as time passes are more realistic and also in line with related research (Detemple et al., 2003; Nielsen and Vassalou, 2006; MacLean et al., 2006; Munk, 2008). MacLean et al. (2006) use a Brownian motion asset pricing model with random parameters and Bayes estimates for the mean return. Nielsen and Vassalou (2006) present a model where first and second moments of stock returns change stochastically over time.

We again assume a risk free asset with zero return and a risky asset with normally distributed return to be available to the investor. The normal distribution of return is described by equation 3.4 in section 3.2.1. As above, we set $\phi = 4$ and cut the tails at $\pm 4\sigma_t$. We again assume a standard deviation of returns of $\sigma_t = 15.0\%$. In contrast to the previous simulation, we now assume μ_t , the mean excess return in period t , to change stochastically over time. Let μ_t be randomly drawn at the beginning of each period t from $\mu_t \sim \mathcal{N}(8.4\%, 3.0\%)$.

In each period we first draw the value for the mean excess return μ_t in that period and then have both investors decide their asset allocations accordingly. The investors have no uncertainty regarding μ_t at time t . Given the values for μ_t and σ_t they can determine their respective allocation of capital to the risky asset in each

¹¹U.S. research returns data (French, 2016).

Table 3.3: Final wealth and bankruptcies of mean-variance investors and growth optimal investors in 1,000 simulation runs with both, time horizons of 100 and 200 periods length and investment opportunities with changing mean returns

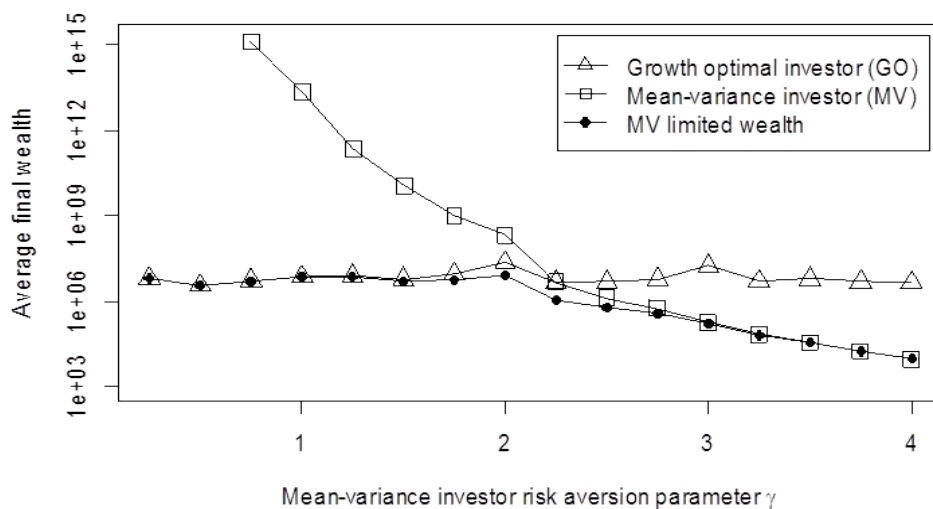
	Investor final wealth over 1,000 simulation runs with ...							
	...100 periods each				...200 periods each			
	Min	Avg	Max	Bankrupt	Min	Avg	Max	Bankrupt
MV $\gamma = 0.5$	NA	NA	2.3e+18	996/1000	NA	NA	NA	1000/1000
MV $\gamma = 1.0$	NA	2.3e+13	1.5e+16	646/1000	NA	7.4e+20	1.6e+23	889/1000
MV $\gamma = 1.5$	NA	1.2e+10	5.4e+12	149/1000	NA	8.2e+17	5.9e+20	249/1000
MV $\gamma = 2.0$	NA	2.2e+8	7.9e+10	31/1000	NA	3.2e+14	2.6e+17	36/1000
MV $\gamma = 2.5$	NA	1.3e+6	1.1e+8	3/1000	NA	1.7e+11	3.9e+13	5/1000
MV $\gamma = 3.0$	91.3	1.8e+5	2.8e+7	0/1000	3.4e+4	5.0e+09	8.6e+11	0/1000
MV $\gamma = 3.5$	18.3	3.5e+4	2.1e+6	0/1000	3.2e+3	2.4e+08	2.2e+10	0/1000
MV $\gamma = 4.0$	39.8	9.4e+3	4.4e+5	0/1000	2.8e+3	1.9e+07	9.2e+8	0/1000
MV lim $\gamma = 0.5$	14.6	3.4e+6	3.2e+8	0/1000	9.3e+2	2.4e+13	1.4e+16	0/1000
MV lim $\gamma = 1.0$	3.4	6.9e+6	1.3e+9	0/1000	1.1e+5	5.4e+12	8.5e+14	0/1000
MV lim $\gamma = 1.5$	11.3	4.7e+6	5.4e+8	0/1000	2.9e+4	2.4e+13	1.2e+16	0/1000
MV lim $\gamma = 2.0$	29.7	8.1e+6	2.8e+9	0/1000	7.3e+5	6.5e+12	4.5e+15	0/1000
MV lim $\gamma = 2.5$	30.4	5.9e+5	5.0e+7	0/1000	1.7e+3	8.7e+10	1.7e+13	0/1000
MV lim $\gamma = 3.0$	83.2	1.6e+5	2.5e+7	0/1000	3.4e+4	4.9e+09	8.2e+11	0/1000
MV lim $\gamma = 3.5$	18.3	3.5e+4	2.1e+6	0/1000	3.2e+3	2.4e+08	2.2e+10	0/1000
MV lim $\gamma = 4.0$	39.8	9.4e+3	4.4e+5	0/1000	2.8e+3	1.9e+07	9.2e+8	0/1000
Growth optimal	15.5	7.5e+6	3.1e+9	0/1000	6.9e+4	1.8e+13	1.4e+16	0/1000

NA refers to bankruptcies or negative values of wealth; MV stands for mean-variance investor; MV lim stands for mean-variance limited investor who never exceeds growth optimal risk taking

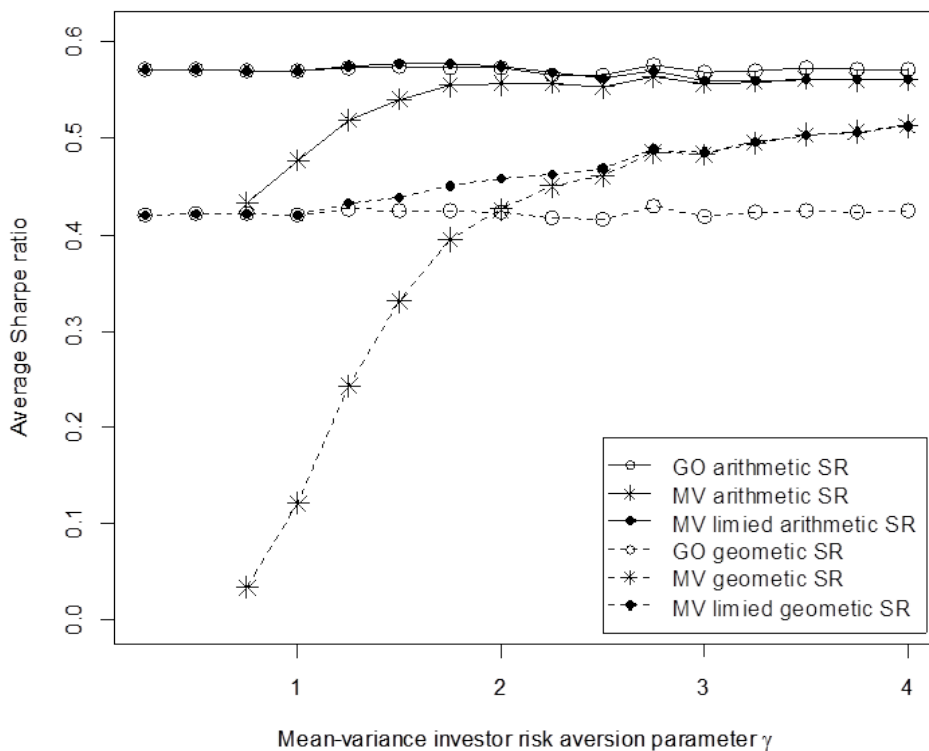
period t . The varying values of μ_t lead our investors to take different asset allocation decisions in each period. In contrast to the previous simulation, where the fraction of wealth that the investors allocated to the risky asset in each period was constant over time, it now changes as μ_t takes different values. The mean-variance investors with different values of risk aversion in this simulation thus do not constantly overbet or underbet, but may oscillate between overbetting and underbetting.

Again, we run 1,000 simulations of 100 periods each for different values of mean-variance investor risk aversion γ . The results of our simulations with stochastic investment opportunities are shown in figure 3.5, as well as in table 3.3 which also contains results for a time horizon of 200 periods. We add one further investor whose investment behavior can be described as follows. At each decision point, the investor looks at the weight in the risky asset that the mean-variance investor and the growth optimal investor use and chooses the lower of these two. We call this investor the mean-variance limited investor, as his risk taking behavior is that of a mean-variance investor who accepts the growth optimal weight in the risky asset as maximum level of risk taking. This investor's behavior fits the suggestion of Markowitz (1976) discussed earlier.

Figure 3.5a shows the average final wealth of our investors. The results for mean-variance and growth optimal investor are similar to those shown in figure 3.4a in



(a) Average final wealth of growth optimal and mean-variance investor



(b) Average arithmetic and geometric Sharpe ratio (SR) for growth optimal investor (GO) and mean-variance investor (MV)

Figure 3.5: Average final wealth (a) and average arithmetic and geometric Sharpe ratio (b) of growth optimal and mean-variance investors with different values of risk aversion $0 < \gamma \leq 4$ over 1,000 simulation runs with 100 periods each and investment opportunities with changing mean returns

section 3.4.1. We note two main differences. First, the mean-variance investor now achieves higher average final wealth than the growth optimal investor for $\gamma > 2.25$. In the previous simulation this was the case for $\gamma > 2$ only. Second, the average final wealth levels of all investors are generally higher, which may be easier seen from table 3.3. The mean-variance limited investor, who never exceeds growth optimal risk taking, either achieves a lower average final wealth than the other two investors, or an average final wealth equivalent to the lower of both other investors.

The number of bankruptcies that each investor experiences over the 1,000 simulations is shown in table 3.3. The mean-variance investors with $\gamma = 1.0, 1.5, 2.0, 2.5$ experience bankruptcy more frequently than in the previous simulation. Also, the mean-variance investor with $\gamma = 2.5$ experiences bankruptcies during the 1,000 simulation runs, although his average final wealth is only about a sixth of that of the growth optimal investor. The mean-variance limited investor does not experience bankruptcies as he does not exhibit overbetting.

Table 3.3 shows again that as the time horizon increases more overbetting mean-variance investors experience bankruptcy. Extending the simulation time horizon to 200 periods we observe the fraction of mean-variance investors who experience bankruptcy to increase for $\gamma = 1.0, 1.5, 2.0, 2.5$. In addition we find the mean-variance investor with risk aversion $\gamma = 0.75$ to achieve an average final wealth below zero while in our simulation with 100 periods in figure 3.5a this investor achieved the highest average final wealth. We find mean-variance investors to overbet more aggressively when investment opportunities are stochastic as they encounter market states with higher mean returns, which leads to ruin more frequently.

Figure 3.5b shows average arithmetic and geometric Sharpe ratios. We see that all depicted mean-variance investors achieve lower average arithmetic Sharpe ratios than the growth optimal investor. In the simulation with constant investment opportunities in section 3.4.1 figure 3.4b we noted that only overbetting mean-variance investors achieve lower average arithmetic Sharpe ratios. It is a new result that the average arithmetic Sharpe ratios are also lower for the non-overbetting mean-variance investors of higher risk aversion. These lower arithmetic Sharpe ratios are explained by the variation in the weight of the risky asset that the stochastic investment opportunities cause. The series of arithmetic returns that mean-variance investors realize tend to be less smooth than that of the growth optimal investor as the mean-variance investor risk taking is more volatile. Figure 3.2 in section 3.3.2.3 visualizes the broader range of weight in the risky asset that the mean-variance investors choose. Our simulation results show that this more volatile risk taking causes the mean-variance investor to experience lower arithmetic returns than the

growth optimal investor, while the standard deviation of his series of arithmetic returns is not correspondingly lower.

The mean-variance limited investor, who never exceeds growth optimal risk taking, will always at least obtain equally high average arithmetic Sharpe ratios as a pure mean-variance investor. We thus conclude that an investor who accepts the arithmetic Sharpe ratio as his primary measure of performance is well advised to consider the growth optimal weight in the risky asset as an upper limit for risk taking in multi-period investment decisions.

In figure 3.5b also the average geometric Sharpe ratio¹² over 1,000 simulations is depicted. Four differences to the previously discussed average arithmetic Sharpe ratio are evident. First, the average geometric Sharpe ratio is generally lower than the average arithmetic Sharpe ratio. This finding and the reasoning is equal to the simulation with constant investment opportunities in section 3.4.1. Second, the performance of overbetting mean-variance optimizers is substantially worse when performance is measured by the geometric Sharpe ratio instead of the arithmetic Sharpe ratio. We have seen and explained this in section 3.4.1 before. Third, mean-variance optimizers with high degrees of risk aversion achieve better geometric mean Sharpe ratios than growth optimal investors. We have also seen and explained this in section 3.4.1. Fourth, and this is the most interesting finding from the simulation with stochastic investment opportunities, we see that the mean-variance limited investor for $1 < \gamma < 3$ outperforms both, the mean-variance investor and the growth optimal investor with regards to average geometric Sharpe ratios. The investment behavior of the mean-variance limited investor does reduce the average final wealth for risk aversion values of $1 < \gamma < 3$ as we see in figure 3.5a, but it results in higher average geometric mean Sharpe ratios. Not only, but especially if an investor accepts the geometric mean Sharpe ratio as his primary measure of performance, he is thus well advised to consider the growth optimal leverage as an upper limit for risk taking in multi-period investment decisions.

3.5 Conclusion

We have presented the risk taking behavior of growth optimal investors and myopic mean-variance investors. The latter keep being used in multi-period settings despite the shortcomings of the myopic assumption. Comparing both investors' choice of weight in the risky asset, we have shown combinations of parameters that result in mean-variance investor overbetting, where overbetting is defined as higher than

¹²Each shown result is an arithmetic average over the 1,000 geometric Sharpe ratios obtained from 1,000 simulation runs.

growth optimal risk taking which for long term investing results in lower wealth accumulation but increased volatility of investor portfolio returns over time. The growth optimal weight in the risky asset can only be determined numerically and there is no analytic solution that allows us to state which parameter combinations result in overbetting. We find that combinations of high risky asset mean return and low standard deviation are particularly prone to overbetting.

We have discussed the effect of overbetting and shown in simulation studies that it can harm both, mean-variance investor wealth accumulation as well as Sharpe ratio. While we find that with finite and short investment horizons the expected final wealth of overbetting mean-variance investors can exceed that of growth optimal investors, these high final wealth levels come at the expense of lower Sharpe ratios and with an increased likelihood of investor bankruptcy. Our simulations illustrate that a series of, from a single period perspective, favorable mean-variance investor asset allocation decisions can lead to almost certain ruin in the long run. In line with the theory of growth optimal strategies we note that the probability of bankruptcy of overbetting mean-variance investors increases with the length of the time horizon.

Our motivation was to follow up on a conjecture made by Markowitz (1976) pointing to the growth optimal leverage as the upper limit for conservative asset allocation in mean-variance analysis when long term wealth growth is desired. Our analysis and simulations have shown that for low values of risk aversion this upper limit to risk taking can indeed be useful to protect mean-variance investor wealth growth and is certainly useful to protect performance.

4 A Performance Comparison of Growth Optimal Strategies and Myopic Mean-Variance Optimization for 80 Years of Stock Market Investing¹

4.1 Introduction

In this paper we apply a growth optimal strategy to stock market investing. We investigate how an investor who follows that strategy would have performed in the stock markets. We compare the growth optimal investor with a myopic mean-variance investor and discuss the benefits and dangers of following either approach. To our knowledge there exists a broad body of research on the theory of growth optimal strategies, but comparatively little on their application to real-world investment situations. We add to the literature by providing novel insights on how growth optimal strategies perform as an asset management strategy. Simultaneously, we highlight the deficiencies of the myopic assumption in multi-period mean-variance optimization when utility functions are not of the isoelastic type.

Growth optimal strategies originated with Kelly (1956) who investigates what bet size a gambler should choose in favorable games with binomial outcomes in order to maximize long-run wealth accumulation. Kelly finds the optimal bet size to be a constant fraction of wealth². More generally, the term growth optimal strategies is used for investment policies that aim at maximizing the expected growth rate of wealth. Kelly (1956) has outlined the intuition of growth optimal strategies in his application of such a strategy to binomial games and following work has established the mathematical foundations and properties of growth optimal strategies. Breiman (1961) analytically proves that the final wealth W_T^* accumulated under a non-terminating growth optimal strategy Λ^* will asymptotically be infinitely larger than any other final wealth W_T that is accumulated under any strategy Λ which is not asymptotically close to the growth optimal strategy Λ^* . Breiman (1961) further finds that the utility function implied by the use of a growth optimal strategy is logarithmic utility. The Kelly criterion and growth optimal strategies are—at least

¹This chapter is based on the homonymous working paper Winselmann (2017b).

²This fraction of wealth is frequently referred to as the Kelly criterion or the full Kelly wager.

for finite time horizons³—from an expected utility theory perspective only optimal for an investor with logarithmic utility (Mossin, 1968; Samuelson, 1969, 1979). In contrast to what some authors initially suggested (Bicksler and Thorp, 1973; Hakansson, 1971b; Rubinstein, 1976), the attractiveness of growth optimal strategies is not rooted in them being a universal concept to describe the risk preferences of all investors or to explain asset prices, but the attractiveness of growth optimal strategies is solely rooted in their superior long term wealth accumulation as well as in their ease of implementation, tractability and increased portfolio stability over time (MacLean et al., 2006; Platen and Rendek, 2011; Ziemba, 2015).

In contrast to growth optimal strategies, modern portfolio theory and mean-variance analysis (Markowitz, 1952, 1959) take into account the different utility functions of individual investors and thus, following the expected utility maxim of von Neumann and Morgenstern (1944), are better suited than growth optimal strategies to solve portfolio selection problems for investors with some defined utility function that is not logarithmic. Nevertheless, modern portfolio theory has originally been developed for a single period framework and its generalization to multi-period investment problems is not straightforward (Markowitz, 1959).

Comparatively easy to solve is the case where a multi-period investment problem is solved optimally by treating each period as if it was the last period. In multi-period problems such behavior of an investor is known as myopic. The conditions for a myopic strategy to be optimal are presented by Mossin (1968) and used by Samuelson (1969) in discrete time and Merton (1969, 1971) in continuous time, but both Samuelson and Merton require restrictive assumptions with regard to utility functions and return distributions.

Indeed, from the time of the fundamental work of Markowitz it took close to half a century until an analytic optimal solution for mean-variance analysis in multi-period portfolio selection was presented by Li and Ng (2000) who derive an analytic optimal portfolio policy and an expression of the mean-variance efficient frontier for multi-period mean-variance analysis using embedding techniques. Their model does not require to restrict utility functions, but it requires the assumption of statistical independence of return vectors over time. Zhu et al. (2004) present an analytic optimal solution which additionally considers risk control over bankruptcy. Wu and Li (2011) follow the same route using embedding techniques and obtain a closed form solution to a multi-period mean-variance portfolio selection problem with regime switching and uncertain exit time. Wu and Li (2012) investigate regime switching

³For infinite time horizons, some authors argue that growth optimal strategies are optimal for any utility function (Markowitz, 1976) while others consider this case irrelevant (Samuelson, 1979).

with an additional stochastic cash flow component and derive a closed form optimal strategy and an efficient frontier. Wu et al. (2014) generalize Wu and Li (2011) to a setting with market state dependent exit time. Cui et al. (2014) amend a no shorting constraint to the setting used by Li and Ng (2000) and derive semi-analytical results which imply a piecewise quadratic value function.

Basak and Chabakauri (2010) present an equilibrium based approach to solve multi-period mean-variance problems. They investigate a dynamic mean-variance problem in a Markovian setting and present an explicit time-consistent solution using dynamic programming and the Hamilton-Jacobi-Bellman equation. Basak and Chabakauri (2010) examine investors with constant risk aversion. Björk et al. (2014) extend their work to investors with state dependant risk aversion.

Despite the advancements following Li and Ng (2000), the mathematically rigorous application of mean-variance analysis to multi-period asset allocation problems remains complex and in both academia and practice mean-variance investors in multi-period asset allocation problems are frequently modeled as myopic (Basak and Chabakauri, 2010). Practical reasons include the difficulty to accurately estimate return and risk for the entire time horizon of the problem at $t = 0$, the complexity and computational intensity of true multi-period models, and the lack of appropriate handling of real-world investor constraints (Kolm et al., 2014). The loss in expected utility that errors originating from additional model complexity and estimates cause may also well outweigh the benefit of a theoretically more rigorous approach (Brandt, 2009; DeMiguel et al., 2009).

Next to exploring the performance of growth optimal strategies for stock market investments in general, our analysis in this paper serves well to illustrate the deficiencies of the myopic assumption by comparing the performance of a myopic mean-variance investor with that of a growth optimal strategy.

Existing work on the application of growth optimal strategies frequently relies on simulation studies that contrast the on average high capital accumulation of growth optimal strategies with the inherent risk of such strategies (Bicksler and Thorp, 1973; Ziemba and Hausch, 1985; Hausch and Ziemba, 1985). Consistent with the theoretical foundations of growth optimal strategies, all studies show high expected growth rates of wealth but also substantial volatility of terminal wealth outcomes and the possibility of high losses for finite time horizons. To reduce risk and volatility, fractional Kelly strategies have been proposed (Clark and Ziemba, 1987; MacLean and Ziemba, 1991; MacLean et al., 1992, 2010a; Ziemba, 2015). Fractional Kelly refers to investing a constant fraction of the growth optimal wager at any decision point, which increases security at the expense of growth. Other approaches to risk

control in the context of growth optimal strategies include drawdown constraints or value at risk (MacLean et al., 2004, 2006), minimum level of wealth constraints (MacLean et al., 2005) or convex shortfall penalties (MacLean et al., 2015). Next to rather theoretically-focused contributions to the field of growth optimal strategies, applications to real market data have been presented. These applications center around two topics. One is gambling, particularly horse betting (Hausch et al., 1981; Thorp, 1984; Hausch and Ziemba, 1985; Bain et al., 2006) and the other is the exploitation of capital market anomalies such as the turn of the year effect (Clark and Ziemba, 1987; Ziemba, 1994, 2012).

MacLean et al. (2011) discuss the performance of growth optimal strategies. While the title of their publication⁴ suggests similarity to our work, MacLean et al. instead focus on extending the simulation-based findings of Bicksler and Thorp (1973) and Ziemba and Hausch (1985). They do not test growth optimal strategies with real world time series data. Additionally, Bicksler and Thorp (1973) and MacLean et al. (2011) do not state exactly how they model their growth optimal investor, but they must be using some approximation of the growth optimal asset allocation that noticeably differs from the actual growth optimal asset allocation, as bankruptcy occurs for a 0.78-fractional Kelly strategy in MacLean et al. (2011). While this is not discussed further in their paper, this contradicts a key property of growth optimal strategies, which is that they never risk ruin (Hakansson, 1971b). In this paper, we use a formulation for the growth optimal investor presented by Rotando and Thorp (1992) which rules out investor ruin when no parameter estimation errors are present.

We extend the existing work on simulation studies of growth optimal strategies in two ways. First, we employ real-world stock market data to assess the performance that a growth optimal investor and a myopic mean-variance investor would have actually achieved. Second, we compare the performance of our growth optimal investor with that of our myopic mean-variance investor. To our knowledge, neither the first nor the second has been exhaustively discussed in the literature yet. We do not only find the wealth accumulation of the growth optimal investor to substantially surpass that of any mean-variance investor, but also find the growth optimal investor and fractional Kelly strategies to achieve superior performance measures.

⁴MacLean et al. (2011) “How does the Fortune’s Formula Kelly capital growth model perform?”

4.2 The model

We model the investment decisions of a growth optimal investor and a mean-variance investor for stock market data from 1926 to 2015. Since continuous re-balancing is impractical, we work in a discrete time setting and our investors take investment decisions in constant time intervals (Rogers, 2001). At each decision point, which is the beginning of each period, the investors re-balance their holdings between two assets. One asset is risk free and the other is risky. Our two-asset setting is thus equivalent to investing along the capital market line (Tobin, 1958). The risk free asset is short term government debt and the risky asset is the market portfolio. We use the stock market aggregate from the U.S. research return data series of French (2016) as a proxy for return on the market portfolio. Following Rotando and Thorp (1992), we assume stock market returns to be normally distributed. Rotando and Thorp acknowledge that the log-normal distribution may provide a superior fit to stock market returns, but given that a log-normal return distribution makes the computations for the growth optimal investor much more cumbersome, they proceed with the assumption of a normal distribution of returns and so do we. In section 4.3.3 we further explore the distribution of returns and confirm that the assumption of normally distributed returns is reasonable.

In the following, we first introduce the myopic mean-variance investor and we second introduce the growth optimal investor. Both investors determine their weight in the risky asset based on the expected value and the standard deviation of the normally distributed excess returns on the risky asset. Return on the risky asset in period t is given by r_t and the certain return on the risk free asset is given by $r_{f,t}$. The expected excess return on the risky asset is given by $E[r_t - r_{f,t}]$ and can equally be written as $E[r_t] - r_{f,t}$ and as $\mu_t - r_{f,t}$. The standard deviation of returns on the risk asset is given by σ_t . The decision variable at the beginning of each period t is the investor's weight ω_t in the risky asset where $\omega_t = 0$ represents 100% of wealth invested in the risk free asset and $\omega_t = 1$ represents 100% of wealth invested in the risky asset. We restrict our investors to not take short positions in the risky asset and thus $\omega_t \geq 0 \forall t$. We assume that the investors can borrow and lend at the risk free rate. Values of $\omega_t > 1$ indicate that an investor takes a short position in the risk free asset and invests more than 100% of wealth in the risky asset. Further, we assume that the investors incur no trading cost.

4.2.1 Mean-variance investor

Preferences for mean-variance investors can either be assumed to be given or they can be derived from expected utility theory. Both approaches are frequently used in the literature (Campbell and Viceira, 2002; Brandt, 2009; Markowitz, 2014). We treat preferences as given and assume a mean-variance investor that is guided by

$$EU_t = E[r_{p,t} - r_{f,t}] - \frac{\gamma}{2}\sigma_{p,t}^2 \quad (4.1)$$

where $r_{p,t}$ is the return on the investor portfolio in period t and $\sigma_{p,t}$ is the standard deviation of the return on the investor portfolio. EU denotes an expected utility function and E is the expectation operator. The risk free rate is given by $r_{f,t}$. Risk aversion is given by γ and is assumed to be constant over time.

The mean-variance investor described in equation 4.1 acts myopic and maximizes his expected utility EU_t in each period t independent of previous and posterior periods. We can equally express $E[r_{p,t} - r_{f,t}]$ as $E[r_{p,t}] - r_{f,t}$ and the certain risk free rate thus does not influence optimal decisions. We use the expected excess return on the risky asset in equation 4.1 for the purpose of more obvious consistency to our formulation of the growth optimal investor presented later.

The investor portfolio return is driven by the weight in the risky asset that the investor chooses. We denote the weight in the risky asset in period t by ω_t and the risky asset return by r_t . The investor portfolio return $r_{p,t}$ thus can be computed as $r_{p,t} = \omega_t r_t + (1 - \omega_t)r_{f,t}$. The investor portfolio excess return is given by $r_{p,t} - r_{f,t} = \omega_t(r_t - r_{f,t})$.

To obtain the expected utility maximizing weight in the risky asset, we first substitute the expected portfolio excess return

$$\begin{aligned} E[r_{p,t} - r_{f,t}] &= E[r_{p,t}] - r_{f,t} \\ &= \omega_t E[r_t] + (1 - \omega_t)r_{f,t} - r_{f,t} \\ &= \omega_t(\mu_t - r_{f,t}) \end{aligned} \quad (4.2)$$

where $\mu_t = E[r_t]$ and we second substitute the portfolio variance

$$\sigma_{p,t}^2 = \omega_t^2 \sigma_t^2 \quad (4.3)$$

into equation 4.1, which yields the optimization problem

$$\max_{\omega_t} \left(\omega_t (\mu_t - r_{f,t}) - \frac{\gamma}{2} \omega_t^2 \sigma_t^2 \right). \quad (4.4)$$

Setting the derivative of equation 4.4 with respect to ω_t equal to zero we obtain the weight in the risky asset ω_t^{mv} which maximizes the expected utility of the mean-variance investor. We use *mv* as a superscript to distinguish the optimal risky asset weights of the mean-variance investor and the growth optimal investor. From equation 4.4 it can readily be seen that the second derivative can only be negative which indicates a maximum. This is a standard textbook approach (Campbell and Viceira, 2002; Alexander, 2008) and the optimal weight in the risky asset ω_t^{mv} for the mean-variance investor is given by

$$\omega_t^{mv} = \frac{\mu_t - r_{f,t}}{\gamma \sigma_t^2}. \quad (4.5)$$

At the beginning of each period t , the optimal weight in the risky asset for the myopic mean-variance investor is thus obtained from the expected excess return on the risky asset $\mu_t - r_{f,t}$ in period t and its standard deviation σ_t .

4.2.2 Growth optimal investor

The strategy pursued by a growth optimal investor is to maximize the expected logarithm of final wealth. In the literature, this is frequently referred to as the ‘Kelly strategy’. Breiman (1961) analytically proves that any essentially different strategy must yield inferior long run wealth growth. From an expected utility theory perspective (von Neumann and Morgenstern, 1944), a growth optimal strategy implies logarithmic utility in wealth. By the properties of the logarithmic utility function, a multi-period problem decomposes into a sequence of independent single period problems and myopic portfolio choice is optimal (Mossin, 1968).

Let W_t be investor wealth at the end of period t . T denotes the time horizon. EU denotes an expected utility function and E is the expectation operator. The investor portfolio return in period t again is given by $r_{p,t}$. Equation 4.6 shows that under logarithmic utility myopic behavior is optimal and that the maximization of the

expected logarithm of final wealth is equivalent to the single-period maximization of the expected log growth rate of wealth.

$$\begin{aligned}
 EU &= E[\log(W_T)] \\
 &= E\left[\log\left(W_0 * \prod_{t=1}^T (1 + r_{p,t})\right)\right] \\
 &= \log(W_0) + \sum_{t=1}^T E[\log(1 + r_{p,t})]
 \end{aligned} \tag{4.6}$$

Rotando and Thorp (1992) have outlined an application of the Kelly criterion (Kelly, 1956) to normally distributed stock market returns. They illustrate their mathematical work with a simple stock market application using yearly S&P500 data from 1928 to 1984. Rotando and Thorp (1992) determine ex-post the constant fraction of wealth that a Kelly investor should have invested into the stock market from 1928 to 1984 in order to maximize the expected logarithm of excess returns. The focus of Rotando and Thorp (1992) is the mathematical groundwork required to generalize the Kelly criterion to continuous distributions of outcomes. Their example merely is a thought experiment, as no investor could in 1928 make investment decisions based on the yet unknown actual parameters of the return distribution of a future time period. We use the model of Rotando and Thorp in a different setting which mirrors the situation of an actual investor who has nothing but past data available to estimate future returns and their volatility. At the beginning of each period t we compute estimates for return and volatility in that period, based on which the investor allocates his capital. Thus, in contrast to Rotando and Thorp (1992), we do not optimize over a past period from an ex-post perspective, but we backtest with historic data an investment policy that a Kelly investor could have actually followed. We deal with the inherent uncertainty of parameter estimates which characterizes real-world investment situations.

In their generalization of the Kelly criterion to normally distributed returns, Rotando and Thorp (1992) assume returns to follow a normal distribution density function on the interval $[A, B]$ with $A = \mu_t - \phi\sigma_t$ as lower limit and $B = \mu_t + \phi\sigma_t$ as upper limit. Rotando and Thorp use $\phi = 3$, which specifies the width of the interval in which realized returns can lie, and work with yearly excess return data and rebalancing. We use various values for ϕ , but assume ϕ to remain constant over time. The tails of the normal distribution have to be cut in order to obtain meaningful results as a positive probability of infinite losses is not only unrealistic for stock

market investments, but also incompatible with a growth optimal strategy. As A converges to minus infinity, the growth optimal wager converges to zero (Rotando and Thorp, 1992). The probability density function $\psi(r_t)$ of returns r_t is defined as

$$\psi(r_t) = \begin{cases} h_t + \frac{1}{\sqrt{2\pi\alpha_t^2}} e^{-(r_t-\mu_t)^2/2\alpha_t^2} & \forall A \leq r_t \leq B \\ 0 & \forall r_t < A \vee r_t > B \end{cases} \quad (4.7)$$

where h_t and α_t are required to preserve the desired properties of the probability density function $\psi(r_t)$ in the presence of cut tails. The parameter h_t is used to ensure that the value of the integral under the probability density function $\psi(r_t)$ is equal to one. The area that we cut in the tails is added back by means of h_t . Cutting the tails not only reduces the area under a normal distribution density function, but it also reduces the observed standard deviation of the stochastic variable. In order to maintain the desired standard deviation of returns σ_t , we employ α_t with $\alpha_t > \sigma_t$ in equation 4.7. The value of α_t which causes the observed standard deviation of returns r_t to be σ_t can only be obtained numerically. See the original paper by Rotando and Thorp (1992) and the application presented in Winselmann (2018) for further details on the derivation of h_t and α_t .

Rotando and Thorp (1992) have their growth optimal investor maximize excess returns and we follow their approach. While they denote excess returns by a single variable s , we instead use $r_t - r_{f,t}$ for excess returns and thus our growth optimal investor's objective function in equation 4.8 looks slightly different but indeed is equivalent.

$$\begin{aligned} G(\omega_t) &= \int_A^B \log(1 + \omega_t(r_t - r_{f,t})) \psi(r_t) dr_t \\ &= \int_A^B \log(1 + \omega_t(r_t - r_{f,t})) \left[h_t + \frac{1}{\sqrt{2\pi\alpha_t^2}} e^{-(r_t-\mu_t)^2/2\alpha_t^2} \right] dr_t. \end{aligned} \quad (4.8)$$

The integration required to set $G'(\omega_t) = 0$ and to determine the growth optimal weight in the risky asset is non-elementary and cannot be done explicitly (Rotando and Thorp, 1992). We numerically optimize $G(\omega_t)$ to obtain the value ω_t^{go} for which $G(\omega_t)$ has its maximum. We call ω_t^{go} the growth optimal weight in the risky asset, where go is used as a superscript to refer to the growth optimal investor. At each decision point, which is the beginning of each period t , we determine the growth optimal weight in the risky asset ω_t^{go} given some estimates for μ_t and σ_t .

4.3 Data and methodology

To evaluate and compare the performance of myopic mean-variance analysis and growth optimal strategies in multi-period investing, we backtest both strategies over an extensive period of stock market returns. We use monthly rebalancing and parameter estimates obtained from 120 past months of stock market excess return data as base case. To test the sensitivity of our results to model parameters, we also present results for different period length and estimation period length as well as for various values of mean-variance investor risk aversion γ and growth optimal investor distribution width ϕ . Lastly, we explore fractional Kelly strategies.

Based on stock market excess return data from 1926 to 2015, we compute the wealth paths which our mean-variance investor and our growth optimal investor would have experienced. In order to obtain a wealth path, we proceed as follows. First, we set the time-invariant parameters. These are the period length, which determines the frequency of re-balancing, and the number of past periods from which estimates for next period returns shall be obtained at each decision point. In addition, we set the time-invariant investor specific parameters ϕ for the growth optimal investor and γ for the mean-variance investor. Second, we compute the estimates for the expected excess return and the volatility of next period returns. Third, we calculate the weights in the risky asset that both investors chose given the estimates for next period expected excess return and volatility. Fourth, we determine investor portfolio returns given the chosen weight in the risky asset and the actual return realization during that period. Steps two to four are repeated for each period. Fifth, we obtain both investors' wealth paths and compute performance measures such as the Sharpe ratio for both investors.

In the remainder of this section we introduce the data set and present the methodology used to obtain estimates. In addition, we take an ex-post perspective and discuss the retrospective quality of our estimates. We find the standardized realized returns to be reasonably close to normally distributed around our estimated mean returns, which supports our assumptions on the distribution of returns and the methodology used to obtain parameter estimates.

4.3.1 Data

We require risk free rate and market return data over a long period. The broader the aggregate market return proxy is, the better it represents a single risky asset. We find the U.S. research returns data from Kenneth R. French's data library (French, 2016) to be the most suitable to our needs. French's monthly data provides the

one-month U.S. Treasury bill rate as a proxy for the risk free rate and an aggregate of several stock indexes as a proxy for the market excess return⁵. The market return proxy used in French's data is broader than in most other data sets and notably broader than any single index which we could alternatively use.

French's data set has monthly return data from mid 1926 to end 2015. Table 4.1 provides descriptive statistics of the data set. Over the entire time period covered by the data set, an imaginary index that started at 1 in mid 1926 would have grown to 4375 by the end of 2015 which is equivalent to an annualized growth of 9.8%. This growth splits into a market excess return effect of factor 216 and a risk free rate effect of factor 20 which represents an annualized growth of 6.2% and 3.4%, respectively. Annual market returns are in the range of -65% to +154%. The one year period with the most severe loss was July 1931 to June 1932 and the one year period with the highest gains directly followed from July 1932 to June 1933.

We do not optimize retrospectively over the entire data set, but instead assume an investor's perspective at each decision point. We assume that at any point in time t we only have knowledge of past market return and risk free rate data until t . Based on this past data, we compute estimates for next period mean excess return and variance which the investors use at the beginning of period t to determine their asset allocation. At the end of period t , we compute the investor portfolio returns from the weight of the risky asset chosen at the beginning of period t and the actual market excess return of that period.

4.3.2 Parameter estimates

At each decision point the investors require estimates for next period returns to decide their asset allocation. We assume return in period t to be normally distributed with mean μ_t and variance σ_t^2 . We follow MacLean et al. (2006) in using moving averages of past return data as estimates for next period mean return and variance. Moving averages are calculated from τ past periods.

$$\begin{aligned}\hat{\mu}_t &= \frac{1}{\tau} \sum_{i=1}^{\tau} r_{t-i} \\ \hat{\sigma}_t^2 &= \frac{1}{\tau-1} \sum_{i=1}^{\tau} (r_{t-i} - \hat{\mu}_t)^2\end{aligned}\tag{4.9}$$

⁵French defines the excess return on the market as the “*value-weight return of all CRSP firms incorporated in the US and listed on the NYSE, AMEX, or NASDAQ that have a CRSP share code of 10 or 11 at the beginning of month t , good shares and price data at the beginning of t , and good return data for t minus the one-month Treasury bill rate (from Ibbotson Associates)*” (French, 2016)

Table 4.1: Descriptive statistics of U.S. research stock returns data set

	Market	Risk free ¹	Market excess ²
Return factor mid 1926 to 2015 ³	4375.1	20.2	216.2
Annualized return	9.8%	3.4%	6.2%
Max annual return	154.5%	15.2%	154.2%
Min annual return	-65.0%	-0.1%	-66.3%
Max monthly return	37.8%	1.4%	37.9%
Min monthly return	-29.1%	-0.1%	-29.1%
Volatility of annual return	25.2%	0.0%	25.4%
Volatility of monthly return	5.4%	0.3%	5.4%

¹ Risk free rate determined as one month U.S. Treasury bill rate; ² Excess return of market over risk free rate; ³ Multiple that a one unit investment in 1926 would have returned end of 2015; Note: The volatility of the risk free rate return refers to the volatility over time. The risk free rate does not have any volatility within each period.

Next to the number of past periods τ from which the moving average estimates are computed, the period length itself requires to be specified. We present results for different period length. Let our base case be a period length of one month and $\tau = 120$. With monthly re-balancing and $\tau = 120$, the moving average estimates $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ are computed from the 120 monthly returns r_{t-120} to r_{t-1} , meaning that the first 10 years of our dataset starting mid 1926 are not available for backtesting as the first month for which we have 120 past months of data available and for which we can obtain estimates is July 1936. We can thus compute estimates for 954 months in the period July 1936 to December 2015 and use this period for backtesting.

4.3.3 Quality of estimates

An important driver of the wealth accumulation and performance of both investors is the accuracy of estimates. Errors in the estimate for the mean return are particularly damaging to wealth growth (Chopra and Ziemba, 1993). Estimation errors in general are known to harm wealth growth significantly more than infrequent portfolio rebalancing (Rogers, 2001).

The less accurate our estimates are, the more far off is the investment decision from what the investors would have chosen if the estimates were accurate. For each pair of mean $\hat{\mu}_t$ and variance $\hat{\sigma}_t^2$ estimates we only have one realization, namely the actual return r_t of that period. We cannot meaningfully compare our estimates with a single return realization, but we can assess the quality of our estimates over the entire time horizon. We compute for each period the distance between the realized return r_t and the mean return estimate $\hat{\mu}_t$ measured in standard deviations $\hat{\sigma}_t$. Given




Figure 4.1: Histogram of deltas between actual return and estimated mean return with normal distribution graph

that we assume normally distributed returns, the normalized deviations should be standard normally distributed. We should observe 5% of realized returns r_t to be further than $\pm 2\hat{\sigma}_t$ from the respective estimated mean $\hat{\mu}_t$ and 68% of realized returns r_t should be within $\pm 1\hat{\sigma}_t$ from the respective mean estimate $\hat{\mu}_t$.

Figure 4.1 shows the distribution of 954 deviations of actual returns from estimated mean returns measured in estimated standard deviations. Period length is one month and 120 past months were used to obtain the estimates, which leaves us with a backtesting period of 954 months or 79.5 years. This is what we above defined as our base case. We observe slight negative skewness with 31% of returns being above $+0.5\sigma$ and while only 28% of returns are below -0.5σ . The highest positive and negative deviations are at $+3.9\sigma$ and -5.3σ , respectively. 95% of all deviations are in the interval $[-2\sigma, +2\sigma]$ and 74% of all deviations are in the interval $[-1\sigma, +1\sigma]$ around the mean. Visual inspection of quantile-normal and normal probability plots support the proximity to normal. We thus conclude that the normal distribution is a reasonably close approximation of the true distribution of returns in our data set and that our assumption of normally distributed returns with estimated mean $\hat{\mu}_t$ and standard deviation $\hat{\sigma}_t$ is justified.

4.4 Results

We evaluate the performance of the myopic mean-variance investor and the growth optimal investor by comparing their wealth paths, their final wealth levels, and their Sharpe ratios. The Sharpe ratio measures the reward-to-variability ratio in the investor wealth path. The use of Sharpe ratios to assess investor performance

in the context of growth optimal strategies is in line with previous literature and particularly multiple publications by Ziemba (Ziemba, 2005; Geyer and Ziemba, 2008; Gergaud and Ziemba, 2012).


We consider an investor bankrupt if his net wealth is below zero at any rebalancing point. In such cases the investor cannot continue to invest and we consider the negative wealth of that period as the final wealth of the investor.

We find the growth optimal investor to achieve higher final wealth levels than the mean-variance investor. This result stands independent of the period length, the number of periods used to obtain estimates, and the value of mean-variance investor risk aversion γ . Our results hold for a wide range of values for ϕ , which for the growth optimal investor specifies the width of the interval on which realized returns can lie. Those values of ϕ for which our results do not hold can be regarded as gross misspecification of the growth optimal investor model. We find that the growth optimal investor not only achieves higher final wealth outcomes, but also achieves higher Sharpe ratios than the mean-variance investor.


4.4.1 Investor wealth paths

In this section we present the wealth paths of two mean-variance investors and two growth optimal investors. We use monthly rebalancing and return data for the time span from mid 1926 to end 2015. Estimates for next period mean return and standard deviation are obtained from 120 months of past data, meaning that investor wealth paths start in mid 1936. The investors we present are first the mean-variance investor with the value of risk aversion γ which yields the highest final wealth of any mean-variance investor and second the growth optimal investor with the value of return distribution width ϕ which yields the highest final wealth of any growth optimal investor. Third and fourth we present a mean-variance investor and a growth optimal investor where γ and ϕ are chosen such that final wealth outcomes are equal. For these four investors we discuss the Sharpe ratios and other performance measures. Figure 4.2 shows the wealth paths of these four investors. Table 4.3 presents performance measures. Our results show that the growth optimal investor performs better than the mean-variance investor by virtually all performance measures.

Figure 4.2a shows the wealth path of a growth optimal investor with $\phi = 7.1$ and a final wealth of \$188,604 at the end of 2015 for each \$1 invested in mid 1936. The growth optimal investor with $\phi = 7.1$ achieves the highest final wealth of any growth optimal investor. Figure 4.2b shows the wealth path of a mean-variance investor with $\gamma = 1.7$ and a final wealth of \$60,880 for each \$1 invested in mid 1936,



(a) Growth optimal investor with $\phi = 7.1$ (b) Mean-variance investor with $\gamma = 1.7$



(c) Growth optimal investor with $\phi = 11.1$ (d) Mean-variance investor with $\gamma = 2.1$

Figure 4.2: Wealth paths of selected investors starting at 1 in mid 1936 with monthly rebalancing and estimation period length of 120 months

which is the highest final wealth of any mean-variance investor. This final wealth is less than one third of what the growth optimal investor with $\phi = 7.1$ achieves.

Figures 4.2c and 4.2d show the wealth paths of a growth optimal investor and a mean-variance investor which both achieve a comparable level of final wealth of \$40,000 for each \$1 invested in mid 1936. The growth optimal investor yields this final wealth level when we set $\phi = 11.1$ and the mean-variance investor yields this final wealth level when we set $\gamma = 2.1$.

The wealth paths in figure 4.2, especially in 4.2c and 4.2d, show that the return generation of the growth optimal investor is more distributed over the entire time span from mid 1936 to end 2015, while the mean-variance investor return generation predominantly occurs in the first half of the time span and particularly from 1950 to 1970. Statistical significance testing confirms this visual finding. The difference between the first half of period returns versus the second half of period returns is statistically significant at the 5% level for the mean-variance investor while it is not statistically significant for the growth optimal investor. The difference between monthly returns from 1950 to 1970 and all monthly returns before and after that

Table 4.2: Differences in returns over time and statistical significance

	Growth optimal investor $\phi = 11.1$		Mean-variance investor $\gamma = 2.1$	
	Average excess return p.a.	t-test for difference	Average excess return p.a.	t-test for difference
First half of sample ¹	15.5%	0.69	21.9%	2.06**
Second half of sample ²	11.4%		7.8%	
1950 to 1970 ³	23.6%	1.77*	36.3%	2.77***
All other months ⁴	9.8%		7.1%	

Notes: Monthly rebalancing, period length 1 month, and estimation period length 120 months

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

¹ 07/1936 to 03/1976, the first 477 months; ² 04/1976 to 12/2015, the second 477 months;

³ 01/1950 to 12/1970, 252 months; ⁴ 07/1936 to 12/1949 and 01/1971 to 12/2015, 702 months

period is statistically significant for both investors, but only at a higher level for the growth optimal investor. Table 4.2 shows detailed results. An investigation of the data shows that between 1950 and 1970 our moving average estimates for stock market excess returns $\hat{\mu}_t$ are high while those for volatility $\hat{\sigma}_t^2$ are relatively low. During this time, both investors weigh the risky asset high in their portfolio, but the mean-variance investor even more so. It plays to the advantage of the mean-variance investor that between 1950 and 1960 the average stock market excess return is particularly high while volatility is rather low, resulting in a steep upward sloping wealth path. After 1960 stock market volatility increases, while our moving average estimates $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ still indicate high returns and relatively low volatility. Given the high weight in the risky asset that the mean-variance investor chooses, his wealth path exhibits high volatility from 1960 to 1970. Winselmann (2018) presents a study on the effects of higher than growth optimal risk taking. Given the harm that excessive risk taking can cause to investor wealth paths, we attribute the high return generation of the mean-variance investor between 1950 and 1970 to a combination of firstly favorable estimates for return distribution parameters that cause his high risk taking and secondly even more favorable actual market developments during that time.

The most commonly used performance measures in literature on growth optimal strategies are the Sharpe ratio and the downside Sharpe ratio (Ziemba, 2005; Geyer and Ziemba, 2008; Gergaud and Ziemba, 2012). We agree and believe that Sharpe Ratio and downside Sharpe Ratio are the most appropriate performance measures. While the downside Sharpe ratio, closely related to the Sortino ratio, seems more appropriate as it does not penalize positive spikes in returns, its lower dispersion in literature and its empirically high correlation with the standard Sharpe ratio (Ziemba, 2005) lets us give preference to the primary use of the standard Sharpe

Table 4.3: Arithmetic and geometric performance measures
 for selected growth optimal and mean-variance investors

	Growth optimal		Mean-variance	
	$\phi = 7.1$	$\phi = 11.1$	$\gamma = 1.7$	$\gamma = 2.1$
Final wealth ¹	188,604	39,989	60,880	39,934
Arithmetic mean excess return p.a. ²	20.4%	13.6%	18.3%	14.8%
Arithmetic standard deviation	39.5%	26.8%	37.8%	30.6%
Sharpe ratio	0.517	0.506	0.484	0.484
Downside Sharpe ratio	0.540	0.528	0.509	0.509
Sortino ratio	0.764	0.746	0.719	0.719
Upside potential ratio	0.633	0.643	0.577	0.577
Calmar ratio ³	0.277	0.289	0.242	0.242
Omega	1.536	1.505	1.562	1.562
Geometric mean excess return p.a. ⁴	12.4%	10.3%	10.8%	10.3%
Geometric standard deviation	54.1%	31.9%	51.9%	38.5%
Sharpe ratio	0.229	0.322	0.208	0.267
Downside Sharpe ratio	0.202	0.306	0.181	0.243
Sortino ratio	0.308	0.453	0.276	0.363
Upside potential ratio	0.424	0.520	0.374	0.421
Calmar ratio ³	0.169	0.219	0.143	0.167
Omega	1.466	1.477	1.443	1.474

Notes: Monthly rebalancing, period length 1 month, and estimation period length 120 months;

¹ Final wealth includes the growth effect of the risk free rate which excess returns exclude;

² Annualized mean excess return over risk free rate; ³ Calculated as annualized mean return divided by maximum monthly drawdown; ⁴ Annualized geometric mean excess return over risk free rate

ratio. In addition, the standard Sharpe ratio can be considered more appropriate in a mean-variance context where also upside variance is perceived as risk. Consistent with Ziemba (2005), we find the Sharpe ratio and the downside Sharpe ratio to yield the same conclusions and similar absolute values, indicating that indeed downside volatility and upside volatility in the investor wealth paths are highly correlated.

Table 4.3 presents performance measures computed from arithmetic and geometric mean and standard deviation. Benchmark return is, where applicable, the risk free rate. For the arithmetic case, we find that both presented growth optimal investors outperform the mean-variance investors with regard to the Sharpe ratio and the downside Sharpe ratio. The growth optimal investors also outperform the mean-variance investors with regard to the Sortino ratio, the upside potential ratio, and the Calmar ratio. The mean-variance investors only outperform when Omega is used to assess performance. The mean-variance investors with $\gamma = 1.7$ and $\gamma = 2.1$

achieve identical values for all performance measures. The reason for this equality is that according to equation 4.5 the risk taking is inversely proportional to the risk aversion parameter γ . With a change in γ we scale the mean-variance investor weight in the risky asset and thus the excess returns of every period with the same factor and consequently observe most performance measures to remain constant.

For the geometric case, we find investment policies which lead to lower final wealth levels to be favored by the performance measures. The $\phi = 7.1$ growth optimal investor outperforms the $\gamma = 1.7$ mean-variance investor with regard to all performance measures and the $\phi = 11.1$ growth optimal investor outperforms the $\gamma = 2.1$ mean-variance investor with regard to all performance measures. While geometric return performance measures tend to favor investment policies that lead to lower final wealth levels, they also favor the growth optimal investor over the mean-variance investor for comparable final wealth levels.

As a robustness check, table 4.A.1 to table 4.A.4 in the appendix break down the entire time span into five sub-periods of equal length. Table 4.A.1 and table 4.A.2 present the same arithmetic performance measures as table 4.3, but for sub-periods of 16 years each. Table 4.A.3 and table 4.A.4 contain the respective geometric performance measures. The results support our findings. Both depicted growth optimal investors achieve consistently superior performance in three out of five sub-periods while in the first sub-period we see a mixed picture and in the third sub-period, in which all investors experience losses, the mean-variance investors achieve mostly better performance measures. In no sub-period we observe both mean-variance investors to outperform both growth optimal investors by all performance measures. For those performance measures and periods where the mean-variance investors are better, we observe smaller differences than where the growth optimal investors are better.

4.4.2 Robustness to variations in ϕ and γ

In this section we analyze the effect of changes in mean-variance investor risk aversion parameter γ and growth optimal investor return distribution width parameter ϕ . As before, we assume monthly rebalancing and compute estimates for next period returns based on 120 past months. Figure 4.3 shows the final wealth outcomes on the primary axis and the corresponding Sharpe ratios computed from the variance in the underlying wealth path on the secondary axis. Final wealth outcomes refer to the wealth level achieved in December 2015 from investing 1 unit in July 1936.

Results for the mean-variance investor are depicted in figure 4.3a. Table 4.4 provides further details. The mean-variance investor achieves a maximum final

wealth of 60,880, which corresponds to a geometric mean annual excess growth rate of 10.8%. The mean-variance investor does not survive for risk aversion $\gamma < 1.3$ but encounters bankruptcy. We define bankruptcy as a negative investor wealth at any rebalancing point. With values of risk aversion $1.3 \leq \gamma < 1.7$ we observe the mean-variance investor to achieve lower final wealth as compared to the less risk averse $\gamma = 1.7$ investor. A detailed study on the reduced wealth accumulation caused by excessive risk taking is provided in Winselmann (2018). For $\gamma > 1.7$ the final wealth decreases and the geometric Sharpe ratios increase. As discussed in section 4.4.1, we find the arithmetic Sharpe ratios of our mean-variance investors to be independent of variations in the risk aversion parameter γ . The Sharpe ratio achieved by the non-bankrupt mean-variance investors with $\gamma \geq 1.3$ is 0.484.

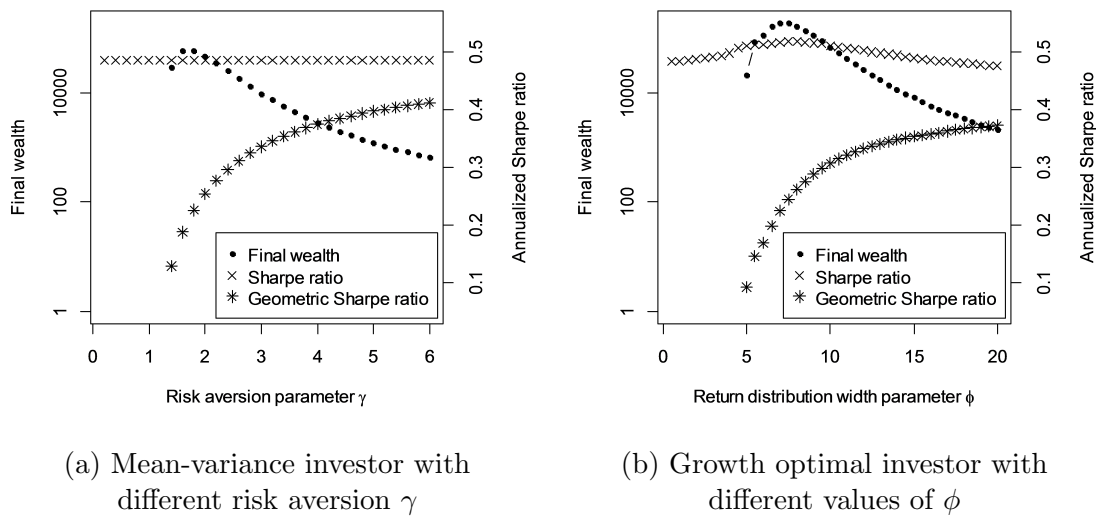


Figure 4.3: Final wealth levels and Sharpe ratios for different values of ϕ and γ with monthly rebalancing and 120 months estimation period

Figure 4.3b shows results for the growth optimal investor. The highest final wealth is achieved with $\phi = 7.1$ at 188,604. Values of $\phi < 5$ cause the growth optimal investor to experience bankruptcy. Such low values constitute a misspecification of the actual return distribution as return realizations to the left of $-5\hat{\sigma}_t$ from $\hat{\mu}_t$ do occur. Values of $5 \leq \phi < 7.1$ result in lower final wealth and lower Sharpe ratios. Values of $\phi > 7.1$ result in lower final wealth, but higher geometric Sharpe ratios. From a reward-to-variability ratio perspective, the lower geometric mean excess returns are thus overcompensated by a decrease in the geometric standard deviation of returns. The maximum arithmetic Sharpe ratio of 0.517 is achieved for $\phi = 7.6$. With regard to the arithmetic Sharpe ratios, the decrease in arithmetic standard deviation is insufficient to compensate for the decrease in arithmetic mean returns for $7.6 < \phi$. Although being less intuitive than its arithmetic counterpart,

Table 4.4: Final wealth levels and corresponding Sharpe ratios of various growth optimal and mean-variance investors

Final wealth ¹	Growth optimal investor			Mean-variance investor		
	Value of ϕ	Sharpe ratio ²	Geometric SR ³	Value of γ	Sharpe ratio ²	Geometric SR ³
200,000	NA ⁴	NA	NA	NA ⁴	NA	NA
188,604	7.1	0.517	0.229	NA	NA	NA
150,000	8.2	0.516	0.267	NA	NA	NA
100,000	9.2	0.514	0.293	NA	NA	NA
60,880	10.2	0.509	0.310	1.7	0.484	0.208
50,000	10.6	0.508	0.316	1.9	0.484	0.241
25,000	12.1	0.501	0.333	2.4	0.484	0.296
10,000	14.3	0.492	0.349	3.0	0.484	0.335
5,000	16.5	0.484	0.360	3.5	0.484	0.358
2,500	19.1	0.477	0.370	4.1	0.484	0.377

Notes: Monthly rebalancing, period length 1 month and estimation period length 120 months; ¹ Assuming initial investment of 1 unit in July 1936; ² Annualized Sharpe ratio; ³ Annualized geometric Sharpe ratio calculated from annualized mean geometric return and annualized geometric standard deviation; ⁴ Wealth level not achieved for any value of ϕ or γ

the geometric Sharpe ratio better depicts the trade-off between risk and volatility in long-term multi-period investment situations. In multi-period problems, arithmetic mean returns are biased upwards and geometric mean returns are more adequate to measure performance (Ziembra, 2005).

Table 4.4 compares final wealth outcomes of growth optimal investors and mean-variance investors in the range of 2,500 to 200,000. If the same level of wealth can be achieved with two different values of γ or ϕ (left and right of the maximum wealth in figure 4.3), the higher value is shown in table 4.4 as the lower value constitutes an inefficient reward-to-variability trade-off. 200,000 is not achieved by any investor. Final wealth above 60,880 is only achieved by growth optimal investors. Table 4.4 shows that the growth optimal investor achieves wealth levels between 5,000 and 60,880 with better Sharpe ratios than the mean-variance investor. Especially with regard to the geometric Sharpe ratio the growth optimal investor outperforms the mean-variance investor by a substantial margin for high levels of final wealth. For a final wealth of 50,000 the respective ratios are 0.316 and 0.241. For a final wealth of 5,000 both investors have the same arithmetic Sharpe ratio while the growth optimal investor still has a slightly better geometric Sharpe ratio. In between 2,500 and 5,000 this changes and the mean-variance investor also achieves better geometric Sharpe ratios. Note that this effect takes place for values of $\phi > 16.5$ and one can discuss how sensible the use of such values is. In section 4.3.3 we have outlined that the worst negative deviation from our moving average return estimates has been a $-5.25\hat{\sigma}_t$ deviation from $\hat{\mu}_t$, which was a -23.2% monthly return in October 1987 when *Black Monday* occurred. It is not a sensible approach to increase the tail length of the distribution in order to decrease risk taking of the growth optimal investor.

Rather, distribution assumptions should be made to best accommodate the actual return distribution and fractional Kelly strategies should be used to accommodate lower investor risk affinity. In section 4.4.5 we discuss fractional Kelly strategies and compare them to mean-variance investors with higher risk aversion. In this section we conclude that our results are fairly robust to variations in ϕ and γ , unless ϕ is increased to heights which we suggest to consider a misspecification of the underlying return distribution.

4.4.3 Robustness to variations in period length

In order to assess the impact of variations in period length we provide results for bi-monthly and trimonthly rebalancing in this section. We keep the estimation period length at 120 months, meaning that mean return and standard deviation estimates are for bimonthly re-balancing obtained from 60 past two-month periods and for trimonthly re-balancing they are obtained from 40 past three-month periods. We observe the same results as in section 4.4.1 and section 4.4.2 above. The growth optimal investor achieves higher final wealth outcomes and better reward-to-variability ratios for high final wealth levels. Tables 4.5 and 4.6 provide a comparison of reward-to-variability ratios for equal final wealth levels. Figure 4.A.1 in the appendix visualizes the results.

For bimonthly rebalancing, the growth optimal investor achieves a maximum final wealth that is more than six times higher than that of the mean-variance investor. Table 4.5 presents the Sharpe ratios and geometric Sharpe ratios for final wealth levels between 500 and 75,000 per 1 unit invested in July 1936. The growth optimal investor achieves higher reward-to-variability ratios for the entire range of values of final wealth. As we increase values of ϕ and γ , we observe the margin to narrow and eventually we would find the mean-variance investor to achieve higher reward-to-variability ratios, but as argued above in section 4.4.2 such high values of ϕ can be considered a misspecification of the underlying return distribution. The appropriate comparison for mean-variance investors with higher levels of risk aversion are the fractional Kelly strategies we investigate in section 4.4.5 below. Figure 4.A.1a and figure 4.A.1b in the appendix depict the results from table 4.5 graphically.

For trimonthly rebalancing we observe a similar picture, but with generally lower final wealth outcomes and reward-to-variability ratios. The growth optimal investor achieves a maximum final wealth that is 4.5 times higher than that of the mean-variance investor. Table 4.6 presents the Sharpe ratios and geometric Sharpe ratios for final wealth levels between 500 and 30,000 per 1 unit invested in July 1936. Figure 4.A.1c and 4.A.1d in the appendix depict the results from table 4.6 graphically.

Table 4.5: Final wealth levels and corresponding Sharpe ratios of growth optimal and mean-variance investors with bimonthly re-balancing and estimates from past 120 months' returns

Final Wealth ¹	Growth optimal investor			Mean-variance investor		
	Value of ϕ	Sharpe ratio ²	Geometric SR ³	Value of γ	Sharpe ratio ²	Geometric SR ³
75,000	NA ⁴	NA	NA	NA ⁴	NA	NA
66,945	6.0	0.496	0.228	NA	NA	NA
50,000	6.7	0.492	0.256	NA	NA	NA
25,000	7.8	0.482	0.283	NA	NA	NA
10,246	9.3	0.471	0.305	2.2	0.444	0.201
5,000	10.8	0.461	0.316	3.0	0.444	0.279
2,500	12.8	0.457	0.335	3.7	0.444	0.314
1,000	16.4	0.452	0.358	4.8	0.444	0.346
500	20.3	0.447	0.372	6.0	0.444	0.367

Notes: Bimonthly rebalancing, period length 2 months and estimation period length 120 months; ¹ For an initial investment of 1 unit in July 1936; ² Annualized Sharpe ratio; ³ Annualized geometric Sharpe ratio calculated from annualized mean geometric return and annualized geometric standard deviation; ⁴ Wealth level not achieved for any value of ϕ or γ

Table 4.6: Final wealth levels and corresponding Sharpe ratios of growth optimal and mean-variance investors with trimonthly re-balancing and estimates from past 120 months' returns

Final Wealth ¹	Growth optimal investor			Mean-variance investor		
	Value of ϕ	Sharpe ratio ²	Geometric SR ³	Value of γ	Sharpe ratio ²	Geometric SR ³
30,000	NA ⁴	NA	NA	NA ⁴	NA	NA
25,137	5.4	0.461	0.207	NA	NA	NA
10,000	7.5	0.456	0.283	NA	NA	NA
5,493	8.6	0.450	0.301	2.3	0.415	0.179
5,000	8.8	0.449	0.304	2.6	0.415	0.216
2,500	10.4	0.444	0.322	3.4	0.415	0.271
1,000	13.1	0.433	0.337	4.6	0.415	0.313
500	16.0	0.422	0.344	5.8	0.415	0.336

Notes: Trimonthly rebalancing, period length 3 months and estimation period length 120 months; ¹ For an initial investment of 1 unit in July 1936; ² Annualized Sharpe ratio; ³ Annualized geometric Sharpe ratio calculated from annualized mean geometric return and annualized geometric standard deviation; ⁴ Wealth level not achieved for any value of ϕ or γ

ically. For trimonthly rebalancing we find similar results as before for bimonthly rebalancing. The growth optimal investor achieves higher reward-to-variability ratios for the entire range of final wealth values. Again, the discussion presented above about higher values of ϕ and γ holds true. Consequently, we conclude that our results are fairly robust to variations in period length and do not depend on monthly re-balancing.

4.4.4 Robustness to variations in estimation period length

We provide this analysis for completeness. We already discussed the quality of our estimates for monthly returns and an estimation period length of 120 months in sec-

tion 4.3.3 above, where we found realized returns to be reasonably close to normally distributed around our estimated mean return $\hat{\mu}_t$ and our estimated variance $\hat{\sigma}_t^2$. This supports the validity of our estimates for next period returns generated from 120 past months of return data. In this section we vary the length of the estimation period. The results should be treated with care and an analysis of the obtained estimates similar to that provided for the case of 120 past months in section 4.3.3 is advisable, but is not provided here due to space limitations and the fact that we have already made our case for 120 months of past data.

Table 4.7: Final wealth and Sharpe ratios of growth optimal and mean-variance investor for monthly re-balancing and various estimation period lengths

Estimation period ¹	Growth optimal investor $\phi = 11.1$			Mean-variance investor $\gamma = 2.1$		
	Final wealth ²	Sharpe ratio ³	Geometric SR ⁴	Final wealth ²	Sharpe ratio ³	Geometric SR ⁴
90 months	53,453	0.505	0.317	8,217	0.432	0.148
100 months	136,298	0.556	0.371	21,512	0.447	0.220
110 months	88,271	0.539	0.355	55,914	0.487	0.272
120 months	39,989	0.506	0.321	39,934	0.484	0.266
130 months	24,458	0.488	0.307	37,742	0.485	0.287
140 months	11,414	0.455	0.276	15,030	0.448	0.253
150 months	12,778	0.465	0.283	9,719	0.434	0.238

Notes: Monthly rebalancing and period length 1 month; ¹ Number of previous months used to obtain estimates for mean and variance of next period returns; ² For an initial investment of 1 unit; ³ Annualized Sharpe ratio; ⁴ Annualized geometric Sharpe ratio calculated from annualized mean geometric return and annualized geometric standard deviation

We find our results for investor performance to be fairly robust to modifications of the estimation period length. Table 4.7 shows final wealth outcomes and Sharpe ratios of two investors with varying estimation period lengths between 90 months and 150 months. As in section 4.4.1 above, the investors shown here are again the growth optimal investor with $\phi = 11.1$ and the mean-variance investor with $\gamma = 2.1$. These investors were selected in section 4.4.1 in order to obtain comparable final wealth outcomes for monthly rebalancing and 120 months estimation period. While we find consistently better Sharpe Ratios for the growth optimal investor, it is not surprising that variation in the estimation period length produce a mixed picture for final wealth outcomes. For an estimation period length of 130 months for example, the mean-variance investor achieves higher final wealth, albeit with respect to the Sharpe ratios the growth optimal investor with $\phi = 11.1$ still performs better. The occasional lead in final wealth that we observe here for the mean-variance investor is mainly caused by the deliberately low risk taking of the growth optimal investor with $\phi = 11.1$, which was chosen in section 4.4.1 in order to obtain equal final wealth outcomes. If we use a growth optimal investor with $\phi = 9.0$ here, we obtain a final wealth of 61,845 and an arithmetic Sharpe ratio of 0.496 which clearly outperforms the mean-variance investor in both dimensions again.

4.4.5 Fractional Kelly strategies

In our robustness tests above we have varied the parameter ϕ and obtained different levels of final wealth for the growth optimal investor, but we should keep in mind that the purpose of ϕ is to correctly specify the width of the distribution of returns expressed in equation 4.7 in section 4.2.2. In contrast to the mean-variance investor risk aversion parameter γ , ϕ is not meant to account for investor risk preferences. While it is important to investigate robustness to changes in ϕ , it would be a mistake to vary ϕ in order to reflect investor risk affinity as this indeed is tampering with the return distribution assumptions.

High values of ϕ do not only reduce growth optimal investor risk taking in general, but in fact they practically also impose a risk taking cap which can be imagined as follows. For simplicity, assume $\hat{\mu}_t = 1\%$ and $\hat{\sigma}_t = 10.1\%$. If we now set $\phi = 10$, this effectively constrains the growth optimal investor's weight in the risky asset $\omega_t^{go} < 1$ as with any higher ω_t^{go} ruin would have a positive probability. With $\phi = 10$, a maximum realized return deviation of ten standard deviations from the mean is permissible, meaning that the minimum possible return realization is $\hat{\mu}_t - \phi\hat{\sigma}_t$ which in this example is $1\% - 10 * 10.1\% = -100\%$. With a positive probability of a total loss, the growth optimal weight in the risky asset is constrained to $\omega_t^{go} < 1$ since it is one of the core properties of growth optimal strategies to never risk ruin (Hakansson, 1971b). Also, in a stock market application it is unrealistic to attach positive probabilities to asset returns of -100% and worse.

While ϕ should be chosen to accurately reflect the distribution of returns, the appropriate means to reduce risk taking rather is to employ fractional Kelly strategies. Fractional Kelly strategies refer to investing a fixed fraction of the growth optimal wager. While the growth optimal investment strategy—in the context of fractional Kelly strategies frequently referred to as *full Kelly*—is optimal in the sense of being utility maximizing for the log utility investor, this is not the case for fractional Kelly strategies. Fractional Kelly strategies are optimal for negative power utility when assets are lognormal (Ziemba, 2015), but they are not fully optimal for any expected utility function in a discrete time settings with normally distributed assets. In such settings, fractional Kelly strategies constitute an approximation and MacLean et al. (1992) argue that fractional Kelly strategies do indeed provide an effective trade-off between growth and security, although not being theoretically optimal.

Table 4.8 presents results for the same levels of wealth as table 4.4 in section 4.4.2 above, but instead of growth optimal investors with varying values of ϕ we here use $\phi = 7.1$ and generate the lower final wealth outcomes by means of fractional Kelly strategies. We find fractional Kelly strategies to yield constant arithmetic

Table 4.8: Final wealth levels and corresponding Sharpe ratios of mean-variance investors with varying risk aversion and fractional Kelly strategies

Final Wealth ¹	Fractional Kelly investor with $\phi = 7.1$			Mean-variance investor		
	Kelly fraction	Sharpe ratio ²	Geometric SR ³	Value of γ	Sharpe ratio ²	Geometric SR ³
200,000	NA ⁴	NA	NA	NA ⁴	NA	NA
188,604	1.00	0.517	0.229	NA	NA	NA
150,000	0.89	0.517	0.264	NA	NA	NA
125,000	0.84	0.517	0.280	NA	NA	NA
100,000	0.90	0.517	0.292	NA	NA	NA
60,880	0.71	0.517	0.319	1.7	0.484	0.208
50,000	0.68	0.517	0.328	1.9	0.484	0.241
25,000	0.59	0.517	0.354	2.4	0.484	0.296
10,000	0.49	0.517	0.383	3.0	0.484	0.335
5,000	0.42	0.517	0.360	3.5	0.484	0.358
2,500	0.36	0.517	0.419	4.1	0.484	0.377

Notes: Monthly rebalancing, period length 1 month and estimation period length 120 months; ¹ Assuming initial investment of 1 unit in July 1936; ² Annualized Sharpe ratio; ³ Annualized geometric Sharpe ratio calculated from annualized mean geometric return and annualized geometric standard deviation; ⁴ Wealth level not achieved for any value of ϕ or γ

Sharpe ratios just as mean-variance investors with different degrees of risk aversion. In the same way as the mean-variance investor risk aversion parameter, the Kelly fraction scales the weight in the risky asset of the growth optimal investor by a constant fraction. This causes mean excess return and standard deviation to have a constant ratio. We find the fractional Kelly investor to be constantly ahead of the mean-variance investor with regard to arithmetic Sharpe ratios. With regard to the geometric Sharpe ratio, we also observe the fractional Kelly investor to achieve better results than the mean-variance investor. In contrast to the results presented in table 4.4 in section 4.4.2 above, the fractional Kelly investor’s geometric Sharpe ratios stay well ahead of those of the mean-variance investor also for lower levels of final wealth. Indeed, for lower levels of final wealth fractional Kelly strategies achieve higher geometric Sharpe ratios than any previously presented investor. Fractional Kelly is not only with regard to our model a more appropriate approach to reduce risk taking, but it also produces superior performance measures as compared to increasing ϕ . These results are in line with MacLean et al. (1992) who found fractional Kelly strategies to provide an effective growth versus security trade-off.


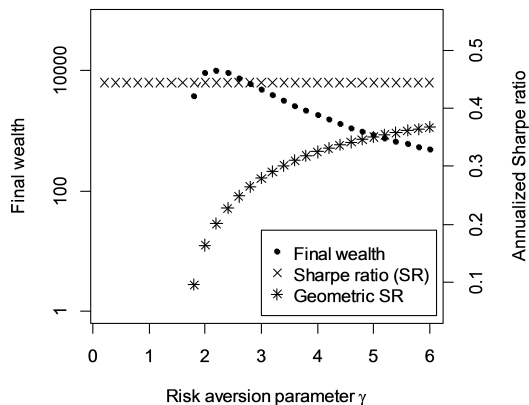
4.5 Conclusion

We investigate a two-assets multi-period capital growth problem in discrete time. The risky asset is represented by the stock market and the return distribution is estimated from past data. We address two issues. First, we explore the performance of a growth optimal investor and a myopic mean-variance investor in a sequence of real world investment situations. Second, we contrast the performance of both

investors and provide evidence on the deficiencies of the myopic assumption that is frequently used in mean-variance optimization. We present final wealth outcomes, wealth paths and performance measures for both investors for 80 years of stock market investing and find the growth optimal investor to achieve higher performance measures for a broad set of parameters. Further, we find fractional Kelly strategies to provide an effective growth versus security trade-off. Model parameters are mean-variance investor risk aversion, width of the return distribution, period length, and number of past periods used to obtain estimates for the distribution of next period returns.


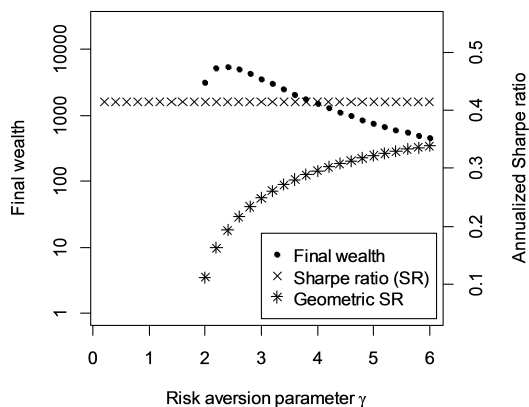
Our results illustrate the deficiencies of the myopic assumption in multi-period mean-variance analysis. Not only does the growth optimal investor accumulate more wealth, which was to be expected given the theory behind it, but he also achieves superior performance measures despite the high volatility in the wealth path which growth optimal strategies are known for. Our results illustrate that the suboptimality of the myopic approach to mean-variance optimization should not be considered a mere theoretical issue, but it indeed matters substantially to investors when multi-period capital growth is desired.

4.A Appendix



(a) Mean-variance investor, bimonthly re-balancing and different risk aversion γ

(b) Growth optimal investor, bimonthly re-balancing and different values of ϕ



(c) Mean-variance investor, trimonthly re-balancing and different risk aversion γ

(d) Growth optimal investor, trimonthly re-balancing and different values of ϕ

Figure 4.A.1: Final wealth levels and Sharpe ratios for different values of ϕ and γ with 120 months estimation period and bi- or trimonthly re-balancing

Table 4.A.1: Arithmetic performance measures of growth optimal investor with $\phi = 7.1$, mean-variance investor with $\gamma = 1.7$ and market in five sub-period of equal length

	Full period 1936-2015			Sub-period 1936-1951			Sub-period 1952-1967			Sub-period 1968-1983			Sub-period 1984-1999			Sub-period 2000-2015		
	GO ¹	MV ²	MKT ³	GO	MV	MKT	GO	MV	MKT	GO	MV	MKT	GO	MV	MKT	GO	MV	MKT
Mean excess return p.a. ⁴	20.4%	18.3%	7.7%	24.8%	19.6%	10.4%	43.4%	51.3%	10.7%	-4.7%	-2.7%	2.0%	35.5%	25.6%	11.6%	3.2%	-2.1%	3.8%
Standard deviation	39.5%	37.8%	15.8%	28.4%	23.3%	18.9%	49.1%	60.4%	11.7%	28.7%	20.0%	16.6%	47.4%	38.5%	15.1%	36.6%	29.8%	15.6%
Sharpe Ratio	0.517	0.485	0.487	0.874	0.839	0.553	0.883	0.851	0.914	-0.163	-0.134	0.118	0.748	0.665	0.767	0.087	-0.072	0.243
Downside Sharpe Ratio	0.540	0.509	0.506	1.081	1.098	0.576	0.988	0.933	1.044	-0.150	-0.129	0.121	0.761	0.656	0.802	0.081	-0.062	0.237
Sortino Ratio	0.764	0.719	0.715	1.528	1.552	0.815	1.397	1.319	1.477	-0.213	-0.183	0.171	1.076	0.927	1.134	0.115	-0.088	0.335
Upside Potential Ratio	0.633	0.577	0.671	0.875	0.849	0.669	0.863	0.831	0.891	0.333	0.332	0.598	0.702	0.614	0.744	0.481	0.365	0.594
Calmar Ratio ⁵	0.277	0.242	0.322	0.911	0.780	0.439	1.117	0.997	1.232	-0.134	-0.129	0.151	0.482	0.337	0.498	0.079	-0.044	0.221
Omega	1.536	1.562	1.445	2.027	2.131	1.547	1.886	1.854	1.926	0.844	0.862	1.091	1.802	1.781	1.794	1.075	0.935	1.196

Monthly rebalancing, period length 1 month, and estimation period length 120 months; ¹Growth optimal investor; ²Mean-variance investor; ³Passive market investment of all capital;

⁴ Annualized mean excess return over risk free rate; ⁵ Calculated as annualized mean return divided by maximum monthly drawdown

Table 4.A.2: Arithmetic performance measures of growth optimal investor $\phi = 11.1$, mean-variance investor with $\gamma = 2.1$ and market in five sub-periods of equal length

	Full period 1936-2015			Sub-period 1936-1951			Sub-period 1952-1967			Sub-period 1968-1983			Sub-period 1984-1999			Sub-period 2000-2015		
	GO ¹	MV ²	MKT ³	GO	MV	MKT	GO	MV	MKT	GO	MV	MKT	GO	MV	MKT	GO	MV	MKT
Mean excess return p.a. ⁴	13.6%	14.8%	7.7%	15.6%	15.9%	10.4%	27.5%	41.6%	10.7%	-3.2%	-2.2%	2.0%	24.7%	20.7%	11.6%	3.2%	-1.7%	3.8%
Standard deviation	26.8%	30.6%	15.8%	19.3%	18.9%	18.9%	31.1%	48.9%	11.7%	22.5%	16.2%	16.6%	31.8%	31.1%	15.1%	25.8%	24.1%	15.6%
Sharpe Ratio	0.506	0.485	0.487	0.811	0.839	0.553	0.884	0.851	0.914	-0.141	-0.134	0.118	0.776	0.665	0.767	0.126	-0.072	0.243
Downside Sharpe Ratio	0.528	0.509	0.506	0.965	1.098	0.576	0.988	0.933	1.044	-0.130	-0.129	0.121	0.806	0.656	0.802	0.120	-0.062	0.237
Sortino Ratio	0.746	0.719	0.715	1.365	1.552	0.815	1.398	1.319	1.477	-0.184	-0.183	0.171	1.140	0.927	1.134	0.170	-0.088	0.335
Upside Potential Ratio	0.643	0.577	0.671	0.853	0.849	0.669	0.863	0.831	0.891	0.357	0.332	0.598	0.732	0.614	0.744	0.515	0.365	0.594
Calmar Ratio ⁵	0.289	0.242	0.322	0.900	0.780	0.439	1.118	0.997	1.232	-0.117	-0.129	0.151	0.527	0.337	0.498	0.126	-0.044	0.221
Omega	1.505	1.562	1.445	1.866	2.131	1.547	1.887	1.854	1.926	0.870	0.862	1.091	1.825	1.781	1.794	1.106	0.935	1.196

Monthly rebalancing, period length 1 month, and estimation period length 120 months; ¹Growth optimal investor; ²Mean-variance optimizer; ³Passive market investment of all capital;

⁴ Annualized mean excess return over risk free rate; ⁵ Calculated as annualized mean return divided by maximum monthly drawdown

Table 4.A.3: Geometric performance measures of growth optimal investor with $\phi = 7.1$, mean-variance investor with $\gamma = 1.7$ and market in five sub-period of equal length

	Full period 1936-2015		Sub-period 1936-1951		Sub-period 1952-1967		Sub-period 1968-1983		Sub-period 1984-1999		Sub-period 2000-2015		
	GO ¹	MV ²	GO	MKT ³	GO	MV	GO	MKT	GO	MV	GO	MV	MKT
Geometric													
mean excess return p.a. ⁴	12.4%	10.8%	22.9%	18.3%	35.7%	36.8%	10.4%	0.6%	23.0%	16.6%	10.9%	-4.0%	2.6%
standard deviation	54.1%	51.9%	32.5%	25.7%	65.4%	89.4%	12.3%	18.0%	35.0%	22.3%	16.7%	47.7%	17.1%
Sharpe Ratio	0.229	0.208	0.704	0.714	0.545	0.412	0.846	0.032	-0.247	-0.206	0.654	-0.083	0.151
Downside Sharpe Ratio	0.202	0.181	0.805	0.860	0.379	0.943	0.943	0.032	-0.204	-0.182	0.650	-0.069	0.136
Sortino Ratio	0.308	0.276	1.183	1.252	0.604	0.598	1.354	0.046	-0.305	-0.267	0.940	-0.105	0.205
Upside Potential Ratio	0.424	0.374	0.742	0.731	0.582	0.648	0.843	0.542	0.251	0.275	0.662	0.354	0.260
Calmar Ratio ⁵	0.169	0.143	0.841	0.750	0.377	0.918	1.207	0.045	-0.248	-0.221	0.470	-0.098	0.150
Omega	1.466	1.443	2.004	2.080	1.534	1.859	1.936	1.095	0.786	0.815	1.786	1.043	1.200

Monthly rebalancing, period length 1 month, and estimation period length 120 months; ¹ Growth optimal investor; ² Mean-variance investor; ³ Passive market investment of all capital;

⁴ Annualized mean excess return over risk free rate; ⁵ Calculated as annualized mean return divided by maximum monthly drawdown

Table 4.A.4: Geometric performance measures of growth optimal investor $\phi = 11.1$, mean-variance investor with $\gamma = 2.1$ and market in five sub-periods of equal length

	Full period 1936-2015		Sub-period 1936-1951		Sub-period 1952-1967		Sub-period 1968-1983		Sub-period 1984-1999		Sub-period 2000-2015		
	GO ¹	MV ²	GO	MKT ³	GO	MV	GO	MKT	GO	MV	GO	MV	MKT
Geometric													
mean excess return p.a. ⁴	10.3%	10.3%	14.7%	15.1%	25.0%	33.2%	10.4%	0.6%	20.9%	15.8%	10.9%	-0.2%	2.6%
standard deviation	31.9%	38.5%	21.2%	20.4%	36.8%	66.0%	12.3%	18.0%	26.1%	17.6%	16.7%	30.4%	17.1%
Sharpe Ratio	0.322	0.267	0.694	0.740	0.680	0.503	0.846	0.032	-0.216	-0.195	0.654	-0.006	0.151
Downside Sharpe Ratio	0.306	0.243	0.788	0.906	0.414	0.707	0.943	0.032	-0.184	-0.175	0.650	-0.006	0.142
Sortino Ratio	0.453	0.363	1.144	1.312	0.604	1.048	1.354	0.046	-0.272	-0.254	0.940	-0.008	0.205
Upside Potential Ratio	0.520	0.421	0.769	0.754	0.582	0.731	0.843	0.542	0.291	0.286	0.662	0.427	0.282
Calmar Ratio ⁵	0.219	0.167	0.846	0.741	0.377	1.020	1.207	0.045	-0.298	-0.204	0.470	-0.008	0.150
Omega	1.477	1.474	1.865	2.090	1.534	1.906	1.936	1.095	0.831	0.823	1.786	1.093	1.200

Monthly rebalancing, period length 1 month, and estimation period length 120 months; ¹ Growth optimal investor; ² Mean-variance optimizer; ³ Passive market investment of all capital;

⁴ Annualized mean excess return over risk free rate; ⁵ Calculated as annualized mean return divided by maximum monthly drawdown

5 Conclusion

In this dissertation we investigate the performance of growth optimal strategies and myopic mean-variance investing in multi-period investment problems. We present a stock market applications and employ both, simulations and real-world time series data. In our simulations we use return distributions which approximately correspond to those observed in stock markets. With this dissertation we contribute to research first by providing insights on the performance of growth optimal strategies in real-world investment situations and second by providing insights on the extent to which the myopic assumption harms wealth accumulation and performance of mean-variance investors in multi-period investment problems.

Following a brief introduction in chapter 1, we introduce the reader to growth optimal strategies in chapter 2 and outline their foundations as well as their utility theory basis. Second, we systematically review past and current literature on the academic dispute over the merits and demerits of growth optimal strategies. When portfolio theory was not yet widely accepted, growth optimal strategies and the Kelly criterion have briefly been considered a premier model of financial markets and a viable alternative to portfolio theory. Growth optimal strategies were subsequently deemed inadequate as an asset pricing model, but various examples of successful investment applications of growth optimal strategies exist. Chapter 2 concludes with presenting an encompassing review and summary of these applications.

In chapter 3, we present a simulation-based study on investor performance in a multi-period investment problem. The setting we employ is a two asset case where at each decision point investors decide the fraction of wealth to invest into a risky asset versus a risk free asset. We model a growth optimal investor and a myopic mean-variance investor for which we discuss the respective choice of weight in the risky asset, the final wealth outcomes and multiple performance measures. Given the complexity of true multi-period formulations of mean-variance optimization, mean-variance optimizers are frequently modeled as myopic in multi-period problems (Basak and Chabakauri, 2010). We investigate the effect of constraining myopic mean-variance optimizers with quadratic preferences to not exceed growth optimal risk taking. We find that this constraint in discrete time multi-period asset allocation problems increases performance and prevents investor bankruptcy.

In chapter 4, we compare the performance of a growth optimal investor and a myopic mean-variance investor for 80 years of stock market investing. We present the wealth paths of these investors had they both invested from 1926 to 2015 according to their respective preferences in two assets, a stock market aggregate and a risk free asset. We present fractional Kelly strategies to compare the results of the myopic mean-variance investor to this derivative of the growth optimal investor. Our results indicate that the growth optimal investor and fractional Kelly investors outperform the myopic mean-variance investor by various performance measures.

The insights presented in this dissertation highlight the applicability of growth optimal investing and fractional Kelly strategies to real-world investment situations. Not only can a well calibrated growth optimal model achieve high returns, but it can also achieve high performance measures.

Our comparison with myopic mean-variance investing—in a setting which violates the conditions outlined by Mossin (1968) and applied by Samuelson (1969) and Merton (1969, 1971)—highlights the deficiencies of the myopic approach in multi-period problems. While the myopic assumption is desirable to reduce the complexity of the underlying problem, it can cause substantial damage to wealth accumulation and performance if the underlying problem is not solved optimally with myopia.

Substantial research effort has been devoted to multi-period mean-variance problems in discrete time and the analytic formulation of the mean-variance efficient frontier for multi-period problems by Li and Ng (2000) in particular has advanced our knowledge and inspired further research. But despite the work of Li and Ng and other researchers that followed, the myopic assumption is still used in both literature and practice. The arguments for its use range from ease of application to estimation errors entailed by more complex models as well as the incompatibility of theoretically optimal models with real world constraints. Further research is needed to shed more light on the deficiencies of the myopic assumption in real world investment situations as well as on the limitations of other restrictive assumptions with regards to utility functions and return distributions that more complex models bring about and let them deviate from real world situations. Additionally, future research could further explore the performance of growth optimal strategies in a broader range of investment situations. Both, an extension of our results to multiple risky assets or risky assets other than the stock market can be made. Advancements in this field would provide investors with valuable knowledge about the deficiencies of practical or realizable investment policies versus theoretically optimal investment policies.

Bibliography

- Alexander, C. (2008). *Market Risk Analysis: Quantitative Methods in Finance*. John Wiley & Sons Ltd., Chichester, England.
- Ali, M. M. (1977). Probability and utility estimates for racetrack bettors. *Journal of Political Economy*, 85(4):803–815.
- Amir, R., Evstigneev, I. V., Hens, T., and Schenk-Hoppé, K. R. (2005). Market selection and survival of investment strategies. *Journal of Mathematical Economics*, 41(1-2):105–122.
- Arrow, K. J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290.
- Aucamp, D. C. (1978). A comment on geometric mean portfolios. *Management Science*, 24(8):859.
- Bain, R. S., Hausch, D. B., and Ziemba, W. T. (2006). An application of expert information to win betting on the Kentucky Derby, 1981–2005. *The European Journal of Finance*, 12(4):283–301.
- Bajeux-Besnainou, I. and Portait, R. (1997). The numeraire portfolio: a new perspective on financial theory. *The European Journal of Finance*, 3(4):291–309.
- Baldeaux, J., Grasselli, M., and Platen, E. (2015). Pricing currency derivatives under the benchmark approach. *Journal of Banking & Finance*, 53:34–48.
- Basak, S. and Chabakauri, G. (2010). Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23(8):2970–3016.
- Basak, S. and Shapiro, A. (2001). Value-at-risk-based risk management: Optimal policies and asset prices. *Review of Financial Studies*, 14(2):371–405.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36.
- Bicksler, J. L. and Thorp, E. O. (1973). The capital growth model: An empirical investigation. *Journal of Financial and Quantitative Analysis*, 8(2):273–287.

- Björk, T., Murgoci, A., and Zhou, X. Y. (2014). Mean-variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance*, 24(1):1–24.
- Brandt, M. W. (2009). Portfolio choice problems. In Ait-Sahalia, Y. and Hansen, L. P., editors, *Handbook of Financial Econometrics: Tools and Techniques*, pages 269–336. North Holland, Amsterdam, The Netherlands.
- Brandt, M. W. and Santa-Clara, P. (2006). Dynamic portfolio selection by augmenting the asset space. *The Journal of Finance*, 61(5):2187–2217.
- Breiman, L. (1961). Optimal gambling systems for favorable games. *Fourth Berkeley Symposium on Probability and Statistics*, 1:65–78.
- Campbell, J. Y. and Viceira, L. M. (2002). *Strategic asset allocation: Portfolio choice for long-term investors*. Oxford University Press, New York.
- Chopra, V. K. and Ziemba, W. T. (1993). The effects of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, 19(2):6–11.
- Clark, R. and Ziemba, W. T. (1987). Playing the turn-of-the-year effect with index futures. *Operations Research*, 35(6):799–813.
- Cui, X., Gao, J., Li, X., and Li, D. (2014). Optimal multi-period mean–variance policy under no-shorting constraint. *European Journal of Operational Research*, 234(2):459–468.
- Davis, M. and Lleo, S. (2008). Risk-sensitive benchmarked asset management. *Quantitative Finance*, 8(4):415–426.
- Davis, M. and Lleo, S. (2010). Fractional kelly strategies for benchmarked asset management. In MacLean, L. C., Thorp, E. O., and Ziemba, W. T., editors, *The Kelly Capital Growth Investment Criterion*, World Scientific Handbook in Financial Economic Series. World Scientific Publishing Co. Pte. Ltd., Singapore.
- Davis, M. and Lleo, S. (2013). Fractional Kelly strategies in continuous time: Recent developments. In MacLean, L. C. and Ziemba, W. T., editors, *Handbook of the Fundamentals of Financial Decision Making*, World Scientific Handbook in Financial Economic Series. World Scientific Publishing Co. Pte. Ltd., Singapore.
- DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953.

- Detemple, J. B., Garcia, R., and Rindisbacher, M. (2003). A Monte Carlo method for optimal portfolios. *The Journal of Finance*, 58(1):401–446.
- Du, K. and Platen, E. (2014). Benchmarked risk minimization. *Mathematical Finance*, forthcoming (advanced online publication).
- Fama, E. F. and MacBeth, J. D. (1974). Long-term growth in a short-term market. *The Journal of Finance*, 29(3):857–885.
- French, K. R. (2016). U.S. research returns data (downloadable files): Fama/French 3 factors [monthly], retrieved February 10, 2016: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.
- Friedmann, J. (1982). Using the Kelly criterion to select optimal Blackjack bets: Unpublished manuscript, Stanford University.
- Gergaud, O. and Ziemba, W. T. (2012). Great investors: Their methods, results, and evaluation. *Journal of Portfolio Management*, 38(4):128–147.
- Geyer, A. and Ziemba, W. T. (2008). The Innovest Austrian pension fund financial planning model InnoALM. *Operations Research*, 56(4):797–810.
- Hakansson, N. H. (1971a). Capital growth and the mean-variance approach to portfolio selection. *The Journal of Financial and Quantitative Analysis*, 6(1):517–557.
- Hakansson, N. H. (1971b). Multi-period mean-variance analysis: Toward a general theory of portfolio choice. *The Journal of Finance*, 26(4):857–884.
- Hakansson, N. H. (1971c). On optimal myopic portfolio policies, with and without serial correlation of yields. *The Journal of Business*, 44(3):p 324–334.
- Hakansson, N. H. and Ziemba, W. T. (1995). Capital growth theory. In Jarrow, R. A., Maksimovic, A., and Ziemba, W. T., editors, *Finance, Handbooks in Operations Research and Management Science*, volume 9, pages 65–86. North Holland.
- Hausch, D. B. and Ziemba, W. T. (1985). Transactions costs, extent of inefficiencies, entries and multiple wagers in a racetrack betting model. *Management Science*, 31(4):381–394.
- Hausch, D. B. and Ziemba, W. T. (1990). Arbitrage strategies for cross-track betting on major horse races. *The Journal of Business*, 63(1):61–78.

- Hausch, D. B. and Ziemba, W. T., editors (1995). *Efficiency of Sports and Lottery Betting Markets*, volume 9 of *Handbooks in Operations Research and Management Science*. Elsevier, Amsterdam.
- Hausch, D. B., Ziemba, W. T., and Rubinstein, M. (1981). Efficiency of the market for racetrack betting. *Management Science*, 27(12):1435–1452.
- Jean, W. H. (1980). The geometric mean and stochastic dominance. *The Journal of Finance*, 35(1):151–158.
- Kallberg, J. G. and Ziemba, W. T. (1983). Comparison of alternative utility functions in portfolio selection problems. *Management Science*, 29(11):1257–1276.
- Kelly, J. L. (1956). A new interpretation of information rate. *Bell Systems Technical Journal*, 35:917–926.
- Kolm, P. N., Tütüncü, R., and Fabozzi, F. J. (2014). 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*, 234(2):356–371.
- Kraus, A. and Litzenberger, R. H. (1975). Market equilibrium in a multiperiod state preference model with logarithmic utility. *The Journal of Finance*, 30(5):1213–1227.
- Lakonishok, J. and Smidt, S. (1988). Are seasonal anomalies real? A ninety-year perspective. *Review of Financial Studies*, 1(4):403–425.
- Latané, H. A. (1959). Criteria for choice among risky ventures. *Journal of Political Economy*, 67(2):144–155.
- Latané, H. A. (1978). The geometric-mean principle revisited: A reply. *Journal of Banking & Finance*, 2(4):395–398.
- Latané, H. A. (1979). The geometric mean criterion continued. *Journal of Banking & Finance*, 3(4):309–311.
- Leibowitz, M. L. and Krasker, W. S. (1988). The persistence of risk: Stocks versus bonds over the long term. *Financial Analysts Journal*, 44(6):40–47.
- Levy, H. and Markowitz, H. M. (1979). Approximating expected utility by a function of mean and variance. *The American Economic Review*, 69(3):308–317.
- Li, D. and Ng, W.-L. (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, 10(3):387–406.

- Lintner, J. (1965). Security prices, risk, and maximal gains from diversification. *The Journal of Finance*, 20(4):587–615.
- Lo, C. F. (2012). The sum and difference of two lognormal random variables. *Journal of Applied Mathematics*, 2012(3):1–13.
- Long, J. (1990). The numeraire portfolio. *Journal of Financial Economics*, 26(1):29–69.
- Luo, Y., Zhu, B., and Tang, Y. (2013). Dynamic optimal capital growth with risk constraints. *Economic Modelling*, 30:586–594.
- Luo, Y., Zhu, B., and Tang, Y. (2014). Simulated annealing algorithm for optimal capital growth. *Physica A: Statistical Mechanics and its Applications*, 408:10–18.
- MacLean, L. C., Sanegre, R., Zhao, Y., and Ziemba, W. T. (2004). Capital growth with security. *Journal of Economic Dynamics and Control*, 28(5):937–954.
- MacLean, L. C., Thorp, E. O., Zhao, Y., and Ziemba, W. T. (2010a). Medium term simulations of the full Kelly and fractional Kelly investment strategies. In MacLean, L. C., Thorp, E. O., and Ziemba, W. T., editors, *The Kelly Capital Growth Investment Criterion*, World Scientific Handbook in Financial Economic Series, pages 543–574. World Scientific Publishing Co. Pte. Ltd., Singapore.
- MacLean, L. C., Thorp, E. O., Zhao, Y., and Ziemba, W. T. (2011). How does the fortune’s formula Kelly capital growth model perform? *Journal of Portfolio Management*, 37(4):96–111.
- MacLean, L. C., Thorp, E. O., and Ziemba, W. T., editors (2010b). *The Kelly Capital Growth Investment Criterion*, volume 3 of *World Scientific Handbook in Financial Economic Series*. World Scientific Publishing Co. Pte. Ltd., Singapore.
- MacLean, L. C., Thorp, E. O., and Ziemba, W. T. (2010c). Long-term capital growth: the good and bad properties of the Kelly and fractional Kelly capital growth criteria. *Quantitative Finance*, 10(7):681–687.
- MacLean, L. C., Zhao, Y., and Ziemba, W. T. (2006). Dynamic portfolio selection with process control. *Journal of Banking & Finance*, 30(2):317–339.
- MacLean, L. C., Zhao, Y., and Ziemba, W. T. (2015). Optimal capital growth with convex shortfall penalties. *Quantitative Finance*, 16(1):1–17.

- MacLean, L. C. and Ziemba, W. T. (1991). Growth-security profiles in capital accumulation under risk. *Annals of Operations Research*, 31(1):501–509.
- MacLean, L. C. and Ziemba, W. T. (1999). Growth versus security tradeoffs in dynamic investment analysis. *Annals of Operations Research*, 85:193–225.
- MacLean, L. C. and Ziemba, W. T. (2006). Capital growth: Theory and practice. In Zenios, S. A. and Ziemba, W. T., editors, *Handbook of Asset and Liability Management*, Handbook of Asset and Liability Management Series, pages 427–473. Elsevier, Amsterdam.
- MacLean, L. C., Ziemba, W. T., and Blazenko, G. (1992). Growth versus security in dynamic investment analysis. *Management Science*, 38(11):1562–1585.
- MacLean, L. C., Ziemba, W. T., and Li, Y. (2005). Time to wealth goals in capital accumulation. *Quantitative Finance*, 5(4):343–355.
- Markowitz, H. M. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.
- Markowitz, H. M. (1959). *Portfolio selection: Efficient diversification of investments*. Wiley, New York.
- Markowitz, H. M. (1976). Investment for the long run: New evidence for an old rule. *The Journal of Finance*, 31(5):1273–1286.
- Markowitz, H. M. (2014). Mean–variance approximations to expected utility. *European Journal of Operational Research*, 234(2):346–355.
- Markowitz, H. M., Reid, D. W., and Tew, B. V. (1994). The value of a blank check. *The Journal of Portfolio Management*, 20(4):82–91.
- Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257.
- Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.
- Merton, R. C. and Samuelson, P. A. (1974). Fallacy of the log-normal approximation to optimal portfolio decision-making over many periods. *Journal of Financial Economics*, 1(1):67–94.
- Michaud, R. O. (1989). The markowitz optimization enigma: Is ‘optimized’ optimal? *Financial Analysts Journal*, 45(1):31–42.

- Mossin, J. (1968). Optimal multiperiod portfolio policies. *The Journal of Business*, 41(2):215–229.
- Munk, C. (2008). Portfolio and consumption choice with stochastic investment opportunities and habit formation in preferences. *Journal of Economic Dynamics and Control*, 32(11):3560–3589.
- Nielsen, L. T. and Vassalou, M. (2006). The instantaneous capital market line. *Economic Theory*, 28(3):651–664.
- Ophir, T. (1978). The geometric-mean principle revisited. *Journal of Banking & Finance*, 2(1):103–107.
- Ophir, T. (1979). The geometric mean principle revisited: A reply to a reply. *Journal of Banking & Finance*, 3(4):301–303.
- Phelps, E. S. (1962). The accumulation of risky capital: A sequential utility analysis. *Econometrica*, 30(4):729–743.
- Platen, E. (2006). A benchmark approach to finance. *Mathematical Finance*, 16(1):131–151.
- Platen, E. and Rendek, R. (2011). Approximating the numéraire portfolio by naive diversification. *Journal of Asset Management*, 13(1):34–50.
- Poundstone, W. (2005). *Fortune’s Formula: The Untold Story of the Scientific Betting System that Beats the Casinos and Wall Street*. Hill and Wang, New York.
- Pulley, L. B. (1983). Mean-variance approximations to expected logarithmic utility. *Operations Research*, 31(4):685–696.
- Ritter, J. R. (1988). The buying and selling behavior of individual investors at the turn of the year. *The Journal of Finance*, 43(3):701–717.
- Rogers, L. (2001). The relaxed investor and parameter uncertainty. *Finance and Stochastics*, 5(2):131–154.
- Roll, R. (1973). Evidence on the “growth-optimum” model. *The Journal of Finance*, 28(3):551–566.
- Rotando, L. M. and Thorp, E. O. (1992). The Kelly criterion and the stock market. *The American Mathematical Monthly*, 99(10):922–931.

- Roy, A. D. (1952). Safety first and the holding of assets. *Econometrica*, 20(3):431–449.
- Rubinstein, M. (1976). The strong case for the generalized logarithmic utility model as the premier model of financial markets. *The Journal of Finance*, 31(2):551–571.
- Rubinstein, M. (1991). Continuously rebalanced investment strategies. *The Journal of Portfolio Management*, 18(1):78–81.
- Rubinstein, M. (2006). *A history of the theory of investments: My annotated bibliography*. Wiley, Hoboken, New Jersey.
- Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246.
- Samuelson, P. A. (1971). The “fallacy” of maximizing the geometric mean in long sequences of investing or gambling. *Proceedings of the National Academy of Sciences of the United States of America*, 68(10):2493–2496.
- Samuelson, P. A. (1979). Why we should not make mean log of wealth big though years to act are long. *Journal of Banking & Finance*, 3(4):305–307.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27(3):379–423.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442.
- Sharpe, W. F. (1966). Mutual fund performance. *The Journal of Business*, 39(1):119–138.
- Sharpe, W. F. (1994). The Sharpe ratio. *The Journal of Portfolio Management*, 21(1):49–58.
- Sinn, H.-W. (2003). Weber’s law and the biological evolution of risk preferences: The selective dominance of the logarithmic utility function, 2002 Geneva risk lecture. *The Geneva Papers on Risk and Insurance Theory*, 28(2):87–100.
- Snyder, W. W. (1978). Horse racing: Testing the efficient markets model. *The Journal of Finance*, 33(4):1109.
- Steinbach, M. C. (2001). Markowitz revisited: Mean-variance models in financial portfolio analysis. *SIAM Review*, 43(1):31–85.

- Sun, W., Fan, A., Chen, L.-W., Schouwenaars, T., and Albota, M. A. (2006). Optimal rebalancing for institutional portfolios. *The Journal of Portfolio Management*, 32(2):33–43.
- Thorp, E. O. (1969). Optimal gambling systems for favorable games. *Review of the International Statistical Institute*, 37(3):273–293.
- Thorp, E. O. (1984). *The Mathematics of Gambling*. Lyle Stuart, Secaucus, NJ.
- Thorp, E. O. (2006). The Kelly criterion in Blackjack, sports betting, and the stock market. In Zenios, S. A. and Ziemba, W. T., editors, *Handbook of Asset and Liability Management*, volume 1 of *Handbook of Asset and Liability Management Series*, pages 385–428. Elsevier, Amsterdam.
- Tobin, J. (1958). Liquidity preference as behavior towards risk. *The Review of Economic Studies*, 25(2):65.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ.
- Winselmann, K. P. (2017a). Growth optimal strategies and the Kelly criterion: a review of theory and applications. Unpublished working paper.
- Winselmann, K. P. (2017b). A performance comparison of growth optimal strategies and myopic mean-variance optimization for 80 years of stock market investing. Unpublished working paper.
- Winselmann, K. P. (2018). Using the Kelly criterion to protect wealth growth in myopic multi-period mean-variance portfolio selection. Unpublished working paper.
- Wu, H. and Li, Z. (2011). Multi-period mean-variance portfolio selection with Markov regime switching and uncertain time-horizon. *Journal of Systems Science and Complexity*, 24(1):140–155.
- Wu, H. and Li, Z. (2012). Multi-period mean–variance portfolio selection with regime switching and a stochastic cash flow. *Insurance: Mathematics and Economics*, 50(3):371–384.
- Wu, H., Zeng, Y., and Yao, H. (2014). Multi-period markowitz’s mean–variance portfolio selection with state-dependent exit probability. *Economic Modelling*, 36:69–78.

- Zhu, S.-S., Li, D., and Wang, S.-Y. (2004). Risk control over bankruptcy in dynamic portfolio selection: A generalized mean-variance formulation. *IEEE Transactions on Automatic Control*, 49(3):447–457.
- Ziemba, R. E. and Ziemba, W. T. (2013). *Investing in the Modern Age*, volume 6 of *World Scientific Series in Finance*. World Scientific Publishing Co. Pte. Ltd., Singapore.
- Ziemba, W. T. (1994). Investing in the turn-of-the-year effect in the us futures markets. *Interfaces*, 24(3):46–61.
- Ziemba, W. T. (2005). The symmetric downside-risk Sharpe ratio. *Journal of Portfolio Management*, 32(1):108–122.
- Ziemba, W. T. (2011). Investing in the turn-of-the-year effect. *Financial Markets and Portfolio Management*, 25(4):455–472.
- Ziemba, W. T. (2012). *Calendar Anomalies and Arbitrage*, volume 2 of *World Scientific Series in Finance*. World Scientific Publishing Co. Pte. Ltd., Singapore.
- Ziemba, W. T. (2015). A response to professor Paul A. Samuelson’s objections to Kelly capital growth investing. *The Journal of Portfolio Management*, 42(1):153–167.
- Ziemba, W. T. and Hausch, D. B. (1985). *Betting at the racetrack*. Dr. Z Investments, Vancouver.