Department of Finance
Copenhagen Business School
Nordea Wealth Management
Master thesis

Authors: Martin Bucht and Daniel Skov Andersen
Supervisor at CBS: Marcel Fischer
Contact person at Nordea: Anders Gadegaard Nielsen
12 May 2017

The Effect of Investment Horizon on Equity Allocation: An Optimal Portfolio Approach

Abstract

This paper examines the effect of investment horizon on the optimal allocation to equities in order to find evidence of time diversification. By using genetic optimization, we create optimal portfolios for different time horizons. Each portfolio has an ideal equity allocation based on a value at risk (VaR) or an expected utility framework. The optimization is based on a dataset covering U.S. real return data for a number of asset classes for the years 1802–2016.

In 1969, Paul Samuelson provided mathematical proof against time diversification that was reliant on three assumptions; that investors exhibit constant relative risk aversion (CRRA), that asset returns are independently and identically distributed (IID) and that wealth is only a function of returns from financial assets. The subsequent time diversification debate has centred on these three assumptions. This paper provides input to the debate by discussing the validity of, as well as relaxing, the three assumptions. The first assumption is relaxed by introducing a VaR framework and other risk preferences than CRRA. The second and third assumption are relaxed when we find evidence of mean reversion in the equity return data and when we introduce a fixed non-financial asset, respectively.

We find solid historical evidence to support the notion that a higher allocation to equities is optimal for agents with longer investment horizons, and that the time diversification effect is present over time. The mean reversion characteristic in our dataset is sufficiently strong to show equity allocation increasing with time horizon irrespective of VaR or utility framework. The introduction of a fixed non-financial asset leads to more aggressive optimal equity allocations, with time diversification still being present.

Keywords: optimal portfolios, time diversification, genetic optimization, mean reversion, value at risk, expected utility, utility functions, risk aversion, equity risk premium, return simulations, autoregressive model, equity allocation, wealth management, modern portfolio theory, Paul Samuelson, constrained optimization, historical returns.

Number of characters: 140,636

Martin Bucht
CPR:

Daniel Skov Andersen
CPR:

Acknowledgements: We wish to thank our supervisor, associate professor Marcel Fischer, for his helpful comments and insight. We also wish to thank Anders Gadegaard Nielsen, who has been our contact person at Nordea Wealth Management, for choosing us to help Nordea with this topic, and for providing us with the dataset.

Table of Contents

List of tables

List of figures

1. Introduction

This paper investigates the effect of investment horizon on equity allocation by creating optimal portfolios using various asset classes. We find solid historical evidence to support the notion that a higher allocation to equities is optimal for investors with longer investment horizons. Our framework is robust to different extensions and the driving characteristics in the data has persisted over time.

Equities are known to provide a higher mean rate of return than bonds, and in order to obtain higher returns the investor has to accept greater risk. However, there is doubt as to whether the return versus risk trade-off for an investor depends on the investment horizon. When looking at long horizons, returns on stocks have been positive, and the standard deviation of annualized returns has decreased.[1] Hence, there is a belief that, when investing for a long horizon, a portfolio consisting of high-risk assets will outperform a portfolio consisting of low-risk assets. It then follows that the optimal portfolio should have a greater allocation to equities, the longer the investment horizon. This is known as time diversification.

The existence of time diversification has been the subject of a longstanding debate that essentially began with a seminal paper by Samuelson (1969). In this paper, Samuelson (1969) disproved time diversification by showing that investors should not change their exposure to risky assets on the basis of time horizon, given three central assumptions. The three assumptions are: (1) that the investor's utility function exhibits constant relative risk aversion (CRRA); (2) that equity returns are independently and identically distributed (IID); and (3) that wealth is only a function of returns from financial assets. The academic debate has centred on these three assumptions. An important paper focusing on the first two assumptions is Kritzman and Rich (1998) who review how the allocation of risky assets depends on risk preferences and the return process of the risky asset. Hanna and Chen (1997) relaxes the third assumption by introducing non-financial assets, specifically measuring what effect this introduction has on the optimal equity allocation. Proponents of time diversification are market practitioners, and scholars such as Siegel (2014), Thorley (1995) as well as Hanna and Chen (1997), whose arguments mainly rest on historical return patterns. Those who believe that time diversification is a fallacy, most notably Samuelson (1963, 1969, 1971, 1989 and 1994) and Kritzman (1992, 1994 and 1998), base their arguments on rigorous mathematical proof combined with perfectly simulated returns, most often IID.

The time diversification debate is interesting since its contents and possible conclusions are highly relevant to financial markets, specifically asset allocation, and ultimately the global economy. At the end of 2016, pension fund asset managers in 22 major pension markets had over 36 trillion USD in assets under management, which corresponded to 62 percent of GDP in these economies.[2] Therefore, it is reasonable to assume that the prevailing investment philosophy at these institutions matter, even to the average citizen. Pension funds take investment horizon into account when allocating their assets, and some, like Sweden's AP3 pension fund, explicitly engage in dynamic asset allocations that incorporate time diversification.[3] Further, Bennyhoff (2008) shows that a number of professional asset managers consider

[1] For historical stock returns and annualized standard deviation, see figure 1 and table 5.
[2] Willis Towers Watson (2017).
[3] OECD (2015).

investment horizon to be a key factor when recommending asset allocations and suggest higher allocation to equities for longer time horizons.[4]

Due to the relevance of time diversification among market practitioners, together with the lack of consensus within the academic world, Nordea's Wealth Management wanted to know more about the topic and asked us if we could write this paper. Specifically, they informed us that they wanted to know if the risk of equities was the same for investors with different investment horizons. And if the risk of equities differed between investment horizons, how would this affect the optimal portfolio for investors with differing risk preferences? After an initial meeting with Nordea, together with an overview of the topic, we became aware of the extensive academic debate and its lack of consensus. Once we understood the difficulty in providing meaningful additions in the form of mathematical proof for or against time diversification, as well as the limited practical use of a purely theoretical application, we decided on applying a relatively practical approach. Our primary goal has been to produce results that are valuable to a market practitioner such as Nordea.

We apply an optimal portfolio approach in order to achieve this goal. Optimal portfolios are created based on a dataset consisting of historical returns for a number of asset classes. The data set was made available to us by Nordea and includes the following U.S. data for the years 1802–2016: annual returns for stocks, Treasury Bills, Treasury Bonds, gold and the consumer price index. Using genetic optimization, we derive optimal portfolios using overlapping time periods, for different investment horizons, in the data set in order to see if the allocation to equity changes. The portfolios are derived using a value at risk (VaR) method as well as an expected utility framework. The VaR framework chooses the allocation to equity in the beginning of the time period that maximizes wealth at the end of the time period, while avoiding returns lower than a defined target, with a specified probability. The expected utility method chooses the start-of-period equity allocation that maximizes the end-of-period utility, which is a function of the end-of-period wealth.

Even though this paper is skewed towards a practitioner approach, it provides input to the academic debate by relaxing or violation the three assumptions by Samuelson (1969). We test and find that there is mean reversion in our dataset of equity real return data, thus violating Samuelson's second assumption. By applying a VaR framework to this dataset, we violate both the first and second assumption. By using the VaR approach, we find that the optimal equity allocation increases with investment horizon. The paper also applies a CRRA utility function in line with Samuelson's first assumption and find that the optimal equity allocation is higher for investors with longer investment horizons. Even when we relax the second assumption by extending the utility function, so that the function exhibits decreasing relative risk aversion (DRRA) and increasing relative risk aversion (IRRA), we find that the optimal allocation to stocks increases with time. The third assumption is violated when we introduce a fixed non-financial asset. The introduction of this asset leads to more aggressive equity allocation, ceteris paribus. The framework still exhibits time diversification with the optimal equity allocation increasing with time. Furthermore, the

[4] Since Bennyhoff (2008) is from a number of years back, it is worth noting that the prevailing philosophy among professional asset managers is the same today. For a more thorough view, see J.P. Morgan (2017).

paper finds that a time diversification strategy is resistant to future decreases in the equity risk premium (ERP), and that the time diversification effect is still prevalent when limiting the analysis to more recent return data. Lastly, anomalies in the form of kinks, displayed in our primary optimal allocation outputs, seem to be randomly distributed among investment horizons and can be explained by a limited number of time periods in which equities greatly underperform relative to Bonds and Bills.

This paper is organized as follows. Section 2 consists of a literature review which contains an overview of Markowitz's modern portfolio theory, a description of Samuelson's three assumptions, a section including fundamental utility theory, and an outline of the relevant academic papers in the time diversification debate. Section 3 contains the theoretical framework where we present our application of autoregressive models (AR), the formal framework and optimization procedure for the VaR and expected utility approach, as well as our use of genetic optimization. Section 4 presents an overview and evaluation of our data set. Section 5 covers the results of our paper, primarily the outputs showing the optimal allocation for the VaR and expected utility framework, but also outputs depicting the effect of introducing a non-financial asset as well as the time diversification effect when using a restricted version of our data set. Section 6 includes the conclusion and discussion in which we validate our results and discuss the validity of Samuelson's three assumptions.

2. Literature review

The time diversification literature treats the question of whether an investor's investment horizon should have an influence on his or her portfolio allocation, specifically with respect to the allocation of risky assets. Time diversification is summed up by Dorsett and Reichstein (1995) in the following two tenets:

- The longer the investment horizon, the greater should the allocation of stocks and other high-return assets be in the portfolio.
- In the long run, one can essentially be certain that a portfolio consisting of high-risk assets will outperform a portfolio consisting of low-risk assets.

In this section, we recap Harry Markowitz's single-period modern portfolio theory and introduce a multi-period setting. We present Paul Samuelson's theoretical proof that investors should not alter their allocation of risky assets as a result of time horizon, and the three assumptions on which this is based. Proceeding from these assumptions we provide an outline of fundamental utility theory due to its central role in this paper. Lastly we present the two strands of literature into which we have chosen to divide the papers, the theorist and the practitioner strand (papers belonging to different sections are not mutually exclusive). The literature that is represented in the theorist section mainly concerns the debate surrounding risk preferences. The theoretical strand is primarily represented by Samuelson (1963, 1969, 1971, 1989 and 1994) and Kritzman (1992, 1994 and 1998). The literature in the practitioner section uses historical data to estimate future returns, and combines this with plausible and practical expected utility functions to create ideal portfolios. The main proponents of this view are Siegel (2014), Thorley (1995), and Hanna and Chen (1997).

2.1 Markowitz's modern portfolio theory

Before we explore time diversification in more detail, we should first establish the subject in the field of finance. One of the most seminal theories in finance is the modern portfolio theory (MTP) by Harry Markowitz.[5]

Markowitz (1952) adopted the rule that an investor considers expected return a desirable thing and variance of return (a proxy for risk) an undesirable thing. Starting from these rules, Markowitz (1952) created a parametric optimization model that could be generally applied while at the same time being simple enough for theoretical analysis and numerical solutions. According to the model, an asset's risk and return should not be evaluated in a vacuum but as part of a portfolio consisting of several different assets.

Formally, a portfolio, p, is said to be mean-variance efficient, i.e. is considered a superior portfolio, if it has a higher expected return than any other portfolio, q – choosing from a universe consisting of the same assets – with the same variance of return, that is,

$$E[R_p] > E[R_q] \; where \; \sigma_p^2 = \sigma_q^2. \tag{1}$$

[5] In order to highlight the importance of MTP, we should add that the Swedish Adademy of Sciences said – when awarding Markowitz the Nobel Price in 1990 – that "Markowitz's work on portfolio theory may be regarded as having established financial micro analysis as a respectable research area in economic analysis".

Or equivalently it can have a lower variance of return for the same level of expected return, assuming the same available set of assets:

$$\sigma_p^2 < \sigma_q^2 \; where \; E[R_p] = E[R_q]. \tag{2}$$

Equation 3 and 4 below depicts expected return and return variance for the portfolio, p, of n assets, and are derived as follows:

$$E[R_p] = w'R, \tag{3}$$

$$\sigma_p^2 = w'\sum w. \tag{4}$$

In the above equations, w' and w represent, respectively, the row and column vectors of the portfolio weights of n assets, while R denotes the vector of expected returns for n assets, and $\sum$ represents the matrix of variances and covariances between the n assets.

Despite its theoretical importance, the MTP by Markowitz (1952) has been criticised in literature. The critique focuses on the assumptions of MTP; of these assumptions, the one that is most relevant to this literature review is the single-period assumption. The single-period assumption in MPT implicitly relies on two points, the first one is that assets are to be held for one period, and the second one is that the optimization procedure only needs to be performed once, since it is based on a single, static estimation of the inputs, i.e. the variables expected return and return variance in equation 3 and 4.[6]

However, in reality most assets are held for multiple periods and investors treat portfolio decisions as repetitious. Detemple (1986) seeks to further understand asset prices by relaxing the assumption that investors observe the state of the economy, i.e. have perfect information. Detemple (1986) points out that in MPT, observed state variables in the economy affect the parameters expected return and return variance in equation 3 and 4. This implies a continuous change in the two main parameters in MTP, and subsequently a deformation of the opportunity set over time.

Brennan (1997) builds on the work by Detemple (1986) by showing that the effect of potential future learning about return and mean variance parameters increases with the degree of uncertainty over the parameter as well as with the investment horizon. The effect of the agent knowing that he or she will be able to learn more about the stated parameters induces the agent to invest more or less in the risky assets. Brennan (1997) shows that this uncertainty has a significant effect on the investment decisions for an investor with a 20-year horizon. Detemple (1986) and Brennan (1997) show that MTP's single-period assumption stands in contrast to the realities of investing.

2.2 Samuelson's three assumptions

When introducing multiple periods, is it then safe to say that the time diversification tenets in the beginning of this section are true? Can we decrease the relatively higher risk inherent in equities by investing over longer time periods?

In one of his fundamental papers, Samuelson (1963) shows that even though risk is reduced by subdividing (a practice familiar to an insurance company) it is not reduced by increasing the number of

[6] Formal framework is the same as in Bianchi et al. (2016) (p. 16).

periods. Samuelson (1963) exemplifies this by telling a story in which he offered one of his colleagues a favourable coin toss. Samuelson (1963) tells his colleague that he will pay $200 in the event the colleague wins, while the colleague only has to pay $100 if Samuelson wins. The colleague answered that he would not accept the gamble since he would feel the loss more than the gain (even though the expected value of the coin toss was positive). However, Samuelson's colleague added that he would take the bet if he were promised 100 such coin tosses. The colleague's supposed rationale being that 100 tosses, contrary to a single toss, would lead to a more favourable outcome on his part. Specifically, the colleague's intuition was based on a loose application of Bernoulli's Law of Large Numbers; he believed that if the number of tosses approached infinity, the probability of winning would essentially be one. Samuelson (1963) clarifies this misinterpretation of Bernoulli's Law of Large Numbers by pointing out the law refers to averages and not sums. As the gamble is repeated, the distribution of potential outcomes spreads. According to Samuelson (1963), if a person would always refuse to take favourable odds on a single gamble, one must rationally refuse to participate in any finite sequence of such gambles.[7]

In one of his seminal papers, Samuelson (1969) focused on the essential dichotomy of the time diversification debate, which Merill and Thorley (1997) phrased as such: "The objective in the time diversification debate is to compare risk at different time horizons (p. 62)."[8] Samuelson (1969) incorporates expected utility theory in order to convincingly show that investors should not change their exposure to risky assets on the basis of time horizon.

Samuelson's conclusions are based on the following three assumptions:

1. Investors have constant relative risk aversion (CRRA), which means that their percentage exposure to risky assets is independent of their wealth.
2. Investment returns are independently and identically distributed (IID), i.e. they follow a random walk and do not exert mean reversion.
3. Wealth is only a function of returns from financial assets; the wealth attributed to returns from human capital and non-financial assets equals to zero.

If these three vital assumptions hold, it follows mathematically that Samuelson's result is true and time diversification is hence a fallacy. Due to this, the majority of the academic stream of literature surrounding time diversification bring up – often in the form of challenging – one, two or all of these three assumptions.

[7] Samuelson (1963) precedes what today is the core of the time diversification debate. The paper is central, primarily in the sense that it is quoted in many other papers in the time diversification literature. This is mainly because a number of scholars mistakenly refer to Samuelson (1963) when they write about the birth of the time diversification debate, while the debate actually stems from Samuelson (1969). Samuelson (1963) is also significant because it in some ways serve as a precursor for Samuelson (1969).
[8] This quote is included to highlight an opaque difference between Samuelson (1963) and Samuelson (1969), i.e. the fact that the time diversification debate is about the relative risks of horizons with different length, as opposed to comparative performance of strategies of the same horizon, but with different investment frequencies.

2.3 Fundamental utility theory

Since expected utility theory is a central component in the time diversification debate – and consequently central to this paper – this section is included to outline the relevant central parts of utility theory.[9]

In one of his classical papers, Bernoulli (1738) stated that "the determination of the value of an item must not be based on its price, but rather on the utility it yields. The price of an item is dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate. Thus, there is no doubt that a gain of one thousand ducats is more significant to a pauper than to a rich man though both gain the same amount (p. 24)." Stated differently, Bernoulli (1738) claims that it is evident that no valid measurement of risk can be acquired without taking its utility into consideration.

Kritzman (1992) expands on Bernoulli's notion of utility, which by today's economists is referred to as diminishing marginal utility. Technically, Bernoulli's notion of utility, which has profound implications for the theory of risk and investment decisions, assumed that the first derivative of utility with respect to wealth is positive and the second derivative with respect to wealth is negative. Formally, this is expressed

$$U'(W) > 0, U''(W) < 0, \tag{5}$$

where U is utility and W is wealth.

From this assumption of diminishing marginal utility, it follows that an individual will reject what is known as a fair game, i.e. a game where the expected value is zero.[10] This phenomenon, what today is referred to as risk aversion, was by Bernoulli (1738) described as "Nature's admonition to avoid the dice altogether".

However, even though Kritzman (1992) concludes that Bernoulli's view of utility is plausible, he states that one cannot conclude that it describes all agents' attitude toward risk. The insights of Bernoulli (1738) have been generalized into a complete theory of risk preferences. Primarily, economists distinguish between those who are risk-averse (will reject a fair game), risk neutral (will be indifferent to accepting a fair game) and risk-loving (will seek out a fair game). Kritzman (1992) highlights the distinction between the absolute amount of an investor's wealth invested in a risky asset versus the percentage of an investor's wealth invested in a risky asset. The amount or percentage of wealth can be a function that is decreasing, staying constant or increasing with wealth.

Table 1 – Risk aversion[11]

	Absolute	Relative
Decreasing	Increase Risky Amount	Increase Risky Percentage
Constant	Maintain Risky Amount	Maintain Risky Percentage
Increasing	Decrease Risky Amount	Decrease Risky Percentage

[9] The content of this section might be seen as less complicated to some readers, but it is nonetheless included due to the fact that the same concepts can be seen as nebulous by other readers.
[10] It is worth pointing out that Samuelson's colleague could be assumed to have a utility function similar to the one described by Bernoulli (1738) since he would not accept a gamble where the expected value was well in the positive.
[11] Table 1 is a reproduction of table II (p. 19) in Kritzman (1992).

Table 1 provides an overview of the risk aversions described by Kritzman (1992). As shown in table 1, a CRRA utility function – the function referred to in Samuelson's first assumption – results in the percentage allocation to equities being independent of wealth.

2.4 The theorist strand

The papers belonging to this strand tend to make use of simulated return data rather than historical. The papers generally dismiss the notion of time diversification. However, the dismissal of time diversification is dependent on Samuelson's first and second assumption. The main focus lies on the importance and viability of Samuelson's first assumption. Specifically, there exists a discussion as to whether CRRA is the most plausible risk preference of the average agent. Samuelson's third assumption, that wealth is only a function of returns from financial assets, are relaxed by some, but most of the authors take it as a necessary assumption to investigate the topic.

How important is the first assumption made by Samuelson (1969), i.e. that investors exhibit CRRA? Would one be able to show that asset allocation is insensitive to investment horizon by only relying on Samuelson's other two assumptions? Readers of Samuelson (1969) got these questions answered as soon as they read the following article in the same journal, which was a companion paper by Merton (1969). In this paper, Merton (1969) derives the optimal equations for a multi-asset problem with a particular case of a two-asset (a risk-free and a risky asset) model while using a CRRA utility function. This results in a confirmation of Samuelson's finding of a constant percentage allocation to the risky assets. However, Merton (1969) extends his analysis by deriving an explicit solution for the case of constant absolute risk aversion (CARA). From this extension, Merton (1969) states the following: "As one becomes wealthier, the proportion of his wealth invested in the risky asset falls (p. 256)." Consequently, as wealth reaches infinity the investor invests all his or her wealth in the risk-free asset. By relaxing the first of Samuelson's three assumptions, Merton (1969) is able to show that the investment horizon, since the expected return is positive, affects the asset allocation.

The proof against time diversification provided by Samuelson (1969) is hence reliant on his first assumption of utility exhibiting CRRA. A natural question is then if there is consensus around whether CRRA best describes the utility function of the typical investor. Merton (1969) states that he finds CARA less plausible from a behavioural perspective than CRRA. Due to this he applies his derivation of the CARA utility function as an extension he does not put much weight into. However, Thorley (1995) claims that a decreasing relative risk aversion (DRRA) utility function best captures rational investor preferences over different investment horizons.

There is also a discussion as to whether the utility function of the average agent is continuous or not. Samuelson (1989) highlights a situation in which the agent has a discontinuous utility function since he or she requires a minimum level of wealth in order to maintain a certain subsistence level. Kritzman (1994) expands on this by explaining that the agent's lifestyle might be drastically affected were he or she to fall below this minimum level of wealth, while further reductions would be less meaningful. Gollier (2001) also finds proof of discontinuity in the form of increased risk aversion with the introduction of liquidity constraints. In some ways similar to a discontinuous utility function, Panyagometh (2011) uses a value at risk (VaR) approach together with bootstrapping simulation from empirical data. Using a VaR approach, his results indicate that the risk of losing wealth decreases as the investment period increases, and that the investor should therefore increase his or her allocation to equities as the time horizon lengthens.

The fact that Samuelson's first assumption is examined and often criticised in the literature has led to authors taking a step back and evaluating the debate itself. Kritzman and Rich (1998) discuss the meaning of risk and point out that even though it is a fact that the probability of loss falls with time for assets with a positive expected return, the question of whether risk increases or decreases with time has to do with the agent's perception of risk. This is due to the fact that the distribution of terminal wealth increases with time, thus increasing the potential size of losses. An investor with a continuous utility function will experience considerable disutility for all levels of such a potential loss, while an investor with a discontinuous utility function (as in the case of VaR as a risk measure) would be less affected once the losses passes a certain level. Kritzman and Rich (1998) state that "the truth is that risk has no universal definition; rather, like beauty, it is in the eye of the beholder. Unfortunately, for many the time diversification debate has degenerated into a referendum on the meaning of risk, which is futile. For others, it is a debate about mathematical truth, which is absurd (p. 67)."[12]

Stangeland and Turtle (1999) review the discussion surrounding utility functions and state that there is little value in debating the merits of time diversification for the typical investor unless an agreed upon definition of the typical investor exists. In order to mitigate this problem, Stangeland and Turtle (1999) apply an extended power utility function which allows them to take preference variability into account.[13] Through their binomial tree simulation, they show that time diversification is affected by the agent's risk preferences as well as the return process of risky assets. Stangeland and Turtle (1999) also claim that apart from risk preference and return process, there are externalities which are essential when it comes to stating the presence of time diversification: (1) the ability of the agent to change his or her work habits, (2) the frequency of required withdrawal from the agent's portfolio, (3) the existence of a non-tradable asset, e.g. human capital, and (4) the potential for changes of the agent's investment knowledge over time.[14]

Kritzman and Rich (1998) provide an extensive overview of how the allocation of risky assets is dependent on risk preference, i.e. utility function, and return process of the risky asset. Kritzman and Rich (1998) show that for an investor with log wealth utility, the indifference between a safe and a risky asset of equal expected utility persists, irrespective of the length of the investment horizon. This result is not limited to agents with log wealth utility functions, but to all agents who have utility functions exhibiting CRRA, given that returns of the risky asset follow a random walk. Furthermore, Kritzman and Rich (1998) look at the square root utility function and the power utility function; they find that the square root function is less risk-averse than log wealth, while the power utility function implies higher risk aversion than log wealth. They also test with quadratic utility as well as combination utility to conclude that both these utility functions result in a non-constant risky asset allocation over time when the risky asset return

[12] A summary of different definitions of risk can be found in exhibit 1 in Kritzman and Rich (1998) (p. 67).
[13] The extended power utility function in question can be found in Stangeland and Turtle (1999) (p. 2).
[14] A more comprehensive description of these factors, specifically with respect to the effect on the weight of equities resulting from a change in one of these factors, ceteris paribus, can be found in table 2 (p. 9) in Stangeland and Turtle (1999).

process exhibits a random walk. In the case of a random walk process, equity allocation decreases with the number of time periods for agents with quadratic utility, while agents with combination utility increases equity allocation with the number of time periods. A summary of these results can be found in table 2.

Table 2 – The impact of preferences and return characteristics on time diversification[15]

Utility Specification	Absolute Risk Aversion	Relative Risk Aversion	Impact of Time on Equity Allocation: Random Walk	Mean Reversion	Mean Aversion
Log Utility $= \ln(Wealth)$	Decreasing	Constant	Hold Constant	Hold Constant	Hold Constant
Square Root Utility $= (Wealth)^{1/2}$	Decreasing	Constant	Hold Constant	Decrease	Increase
Power Utility $= -1/Wealth$	Decreasing	Constant	Hold Constant	Increase	Decrease
Quadratic Utility $= 25 \times Wealth + 0.1 \times Wealth^2$	Increasing	Increasing	Decrease	Decrease	Decrease
Combination Utility $= \frac{1}{Wealth} + \ln(Wealth)$	Decreasing	Decreasing	Increase	Decrease	Increase

Other scholars and other papers, e.g. Kritzman (1994), Levy and Spector (1996), Van Eaton and Conover (1998) as well as Samuelson (1971 and 1994) acknowledge that the optimal allocation of the risky asset over time depends on the utility function, and that their findings might not hold with alternative utility functions. As early as the 1950s, Roy (1952) pointed to this weakness of the utility function approach, saying that "in calling in a utility function to our aid, an appearance of generality is achieved at the cost of a loss of practical significance and applicability in our results (p. 433)".

Views and criticisms surrounding Samuelson's first and to a lesser extent second assumption are found in several papers in the theoretical strand of literature. The third assumption, that wealth is a function only of returns from financial assets, is discussed less often. However, one notable exception to this is Pástor and Stambaugh (2012) who extend their formal framework to include labour income. Also, Bianchi, Drew and Walk (2014) highlight the fact that income from human capital and other factors of potential income are somewhat overlooked. Campbell and Viceira (2002) point out that it is unreasonable to assume that investors' wealth only consists of financial assets, and that one should optimally take into account human capital as well as non-financial assets. The third externality highlighted by Stangeland and Turtle (1999), the existence of a non-tradable asset such as human capital, also challenges Samuelson's third assumption.

2.5 The practitioner strand

The papers in this strand tend to use historical data to help model future returns. The papers generally find evidence of time diversification. Several authors find mean reversion characteristics in historical data,

[15] Table 3 is a reproduction of exhibit 2 (p. 68) in Kritzman and Rich (1998).

thus violating Samuelson's second assumption of IDD returns. Authors who create optimal portfolios tend to assume CRRA utility, mainly due to plausibility and convenience. One paper relaxes Samuelson's third assumption. Our paper is more in line with the papers in this strand, hence the focus is more on the results of the papers rather than the overarching debate.

One of the main advocates for a greater equity allocation for longer investment horizons is Siegel (2014).[16] He states that when one looks at historical returns, it is clear that stocks have, on average, produced positive real returns in excess of both bonds and bills. Siegel (2014) finds, when looking at return data dating back to year 1802, that stocks, in contrast to bonds and bills, have never produced a negative real return for investment horizons lasting 17 years or more. Siegel (2014) also highlights the fact that bonds and bills are subject to inflation uncertainty. Since the introduction of fiat money without a conversion feature in combination with governments engaging in expansionary monetary policy, fixed income is no longer synonymous with fixed purchasing power. This is echoed by Bodie (1995) who points out that investing in bonds exposes the investor to inflation risk, as well as Cagan (1974) who finds that stocks have historically performed well relative to inflation. Siegel (2014) points to the fact that Samuelson's second assumption, that risky asset returns are IID, does not hold for equities when observing historical data. Rather than being IID, equity returns have according to Siegel (2014) been characterized by mean reversion.

Unlike Siegel (2014) who uses data from one country, Blanchett et al. (2013) use real returns of cash, bonds and stocks for 20 different countries in the DMS dataset which stretches from 1900 to 2012. Due to this Blanchett et al. (2013) are able to use 2,260 years of return data. Blanchett et al. (2013) find mean reversion in the data, thus violating Samuelson's second assumption. They use optimal portfolios and find that the equity allocation tends to increase with investment horizon, and that time diversification does exist, or has at least existed historically. Blanchett et al. (2013) also find that optimal equity allocations for the United States (U.S.) have been relatively similar to the optimal equity allocations for the 20-country average. Another scholar who uses data from outside the U.S. is Panyagometh (2011) who runs his analysis using data from financial markets in Thailand together with a bootstrapping simulation.

Butler and Domian (1991) use past return data (Ibbotson-Sinquefield stock and long-term Treasury bond indexes over the period 1926-1988) with an empirical resampling procedure to estimate the prevalence of time diversification in portfolio risk. Butler and Domian (1991) point out that estimation based on historical data is disadvantaged by the fact that the number of observations, at least in their case, is limited. They also admit that returns and risk premiums do not appear to be constant over time, hence extrapolations from the data must be analysed in combination with knowledge of current market conditions. However, despite these factors, Butler and Domian (1991) claim that historical asset returns represent an objective foundation on which to base forecasts on future investment performance, especially forecasts for longer horizons. Using the past data at their disposal, Butler and Domian (1991) find that the chance of earning less on the stock index than the long-term Treasury bond index tend to fall

[16] The references in this paper refers to the fifth edition of Stocks for the Long Run, however Jeremy Siegel has pushed the same pro time diversification thesis since the first edition of the same book was published in 1994.

with the investment horizon. Specifically, they find that there is a 5% chance of earning less on the stock index than the long-term Treasury bond index when investing over a 20-year horizon.

Barberis (2000) examines how investors with long horizons should alter their allocation to risky assets when predictability in asset returns is introduced. Barberis (2000) finds that even after taking uncertainty about model parameters, mean return and return variance, into account, there is still enough predictability in asset returns to make long-horizon investors allocate more to equities. According to Barberis (2000), all but the most risk-averse investors should allocate 100 percent of their optimal portfolio to equities if their investment horizon is 10 years or more. Barberis (2000) use data spanning from June 1952 through December 1995; a value-weighted index of stocks in combination with returns from U.S. Treasury bills.

Hansson and Persson (2000) use a bootstrap approach to investigate whether the weights for bills and equities vary with time horizon in an optimal portfolio. They use monthly data for US stocks and US Treasury bills from 1900 to 1997. Hansson and Persson (2000) conclude that their evidence supports the existence of time diversification, the equity weights in their efficient portfolios are slightly higher for longer than for shorter horizons.

Dolvin et al. (2010) suggest that financial planners consider a 100 percent allocation to equities for their clients for investment horizons stretching 10 years to more. When less than 10 years is left until retirement a more conservative allocation is preferred. Dolin et al. (2010) use annual returns starting in 1926 and then apply a simulation method to review potential future results.

Lebowitz and Kogelman (1991) focus on the balance between equity, which they use as a proxy for all risky assets, and risk-free assets. They measure downside risk by fulfilling what they refer to as a shortfall constraint. This shortfall constraint is similar to a discontinuous utility function and a VaR approach, e.g. there must be a probability of 5% or less that returns fall below 2% over a one-year horizon. This shortfall constraint is then related to a minimum return threshold. Lebowitz and Kogelman (1991) then combine the allowed probability and the corresponding minimum return with investment horizon. The result of their model shows them that when investing for a relatively shorter horizon, the volatility inherent in equities creates a high probability of poor returns. However, when they increase the investment horizon they find that there is sufficient time to capture the equity risk premium. Hence, Lebowitz and Kogelman (1991) find that the allocation to risky assets should increase with time horizon.

Thorley (1995) is one of few authors in this strand to bring up Samuelson's third assumption. Thorley (1995) uses historical S&P 500 return data from 1947 to 1993 (Ibbotson Associates data) in combination with a CRRA utility function to find that the historical mean reversion phenomenon ends at a four-year horizon. Thorley (1995) points out that these results stand in contrast to conventional wisdom of mean reversion of returns for very long horizons such as 20 years. Thorley (1995) finds the most likely explanation to the break after four years to be the fact that Samuelson's third assumption is unrealistic. According to Thorley (1995), expected utility theory, if applied only to the agent's financial assets, may not be a very useful model of investment behaviour and preferences.

Hanna and Chen (1997) are the only authors who relax Samuelson's third assumption by adding a fixed non-financial asset to their optimization framework. They find that the inclusion of a fixed non-financial asset leads to a higher equity allocation. Hanna and Chen (1997) specifically examines risk tolerance and its implications for investment portfolios. They conclude that even investors who possess a relatively high level of risk aversion should have portfolios with a large allocation to risky assets if their investment horizon is 20 years or more. Hanna and Chen (1997) make use of data obtained from Ibbotson Associates (1996), with real US returns from 1926 to 1995 for six categories of financial assets. In line with the conclusions by Siegel (2014), they find mean-reversion in equity returns, hence Samuelson's second assumption is violated. They then calculate expected utility of all possible portfolios, using overlapping periods within the data set, for different investment horizons, and find that equity allocation is positively related to investment horizon.

3. Theoretical framework

In this section, we explain and justify our choice of theoretical framework, and how the approach relates to a market practitioner such as Nordea. The section also touches on why the approach is important from an academic viewpoint, specifically with respect to Samuelson's three assumptions.

Our approach consists of creating optimal portfolios for different investment horizons using a value at risk (VaR) and an expected utility framework in combination with historical return data. We have chosen to use historical data, instead of simulated returns that are perfectly random, to ensure that our conclusions capture all the information inherent in the historical data set. The section also present our use of autoregressive models and genetic algorithms.

3.1 Autoregressive model and information criteria

The autoregressive (AR) model is a specific application of the linear regression model. It is mainly used to forecast future values of a single time series by regressing the time series value on its past values.[17] Applied to stock returns, the p^{th}-order AR model takes the form of the following formal regression:

$$R_t = \sum_{i=1}^{p} \beta_i \cdot R_{t-i} + e_t, \tag{6}$$

where R_t denotes the stock returns at time t, $\beta_i{}'s$ are estimated coefficients, and e_t is the error term.

We apply the model in equation 6 to stock returns in order to find evidence of historical mean reversion (alternatively momentum or neither). Negative estimated coefficients will indicate mean-reversion because it means that stock returns are negatively related to past returns. If estimated coefficients are positive, it is a sign of momentum due to the fact that stock returns are positively related to past values. If mean reversion is found, it adds to the argument for time diversification due to low returns being followed by high returns and vice versa. Essentially, this tells us that the risk of stocks decreases over time since large deviations from the average return are relatively unlikely.

When deciding on p, i.e. the number of time-lagged values to be included, we make use of an information criterion, specifically the Akaike information criteria (AIC). The AIC chooses the order p through a trade-off between the marginal benefit of including more lags to increase the goodness of fit and the marginal cost of the additional complexity added to the model. Having too few lagged variables increases the risk of missing important information contained in the omitted variables, and having too many increases the level of over-fitting, thus increasing the estimation error.[18]

Formally, the AIC is defined as follows:

$$AIC(p) = ln\left(\frac{SSR(p)}{T}\right) + (p+1) * \frac{2}{T}. \tag{7}$$

[17] Stock & Watson (2012) (p. 571).
[18] Stock & Watson (2012) (p. 584-586).

In equation 7 the $SSR(p)$ is the sum of squared residuals of the estimated $AR(p)$, T is the total number of periods, and p is the number of lags used. The AIC gives us different values for different number of lags, the lower the AIC value the better is the specification in question according to the AIC.[19]

In our application, we incorporate the recommendation of the AIC in making our decision regarding the amounts of lags in our AR model; this is mainly because we believe using the AIC constitutes a less parsimonious approach than simply relying our own intuition.

3.2 Value at risk

Value at risk (VaR) is a measure of the risk of investments. In this paper, we will focus on percentage VaR, which will hence be what we refer to when we write VaR. Simplified, VaR summarizes the worst percentage loss over a target horizon that will not be exceeded with a specific certainty. For example, a portfolio might have a one-year VaR of -10 percent at a 95 percent confidence level, this means that for an investment period of one year we can say with 95 percent certainty that the portfolio will not lose more than 10 percent of its initial value.[20]

Similar to Panyagometh (2011), we use VaR as a risk measure in order to construct optimal portfolios. The main reason is because VaR is used by market practitioners such as Nordea's Wealth Management. In turn, this is due to its pedagogic features, i.e. when using VaR it is relatively easy to explain the risk of a portfolio to a layman investor. VaR is also increasingly used since it is a risk measure that provides for an aggregate view of a portfolio's risk. As a result of its usability and relative simplicity, regulators demand its use for better control of financial risks.[21] Furthermore, regulators believe a risk measurement that is relatively easy to understand will serve to mitigate the degree of asymmetric information between the bank and the client in financial transactions.[22] The general mathematical definition of VaR is the following (not corresponding to percentage VaR but depicts VaR for real values):[23]

$$VaR_{\gamma}(W) = inf\{x \in \mathbb{R}: P(W + x < 0) \leq 1 - \gamma\}. \tag{8}$$

In equation 8, VaR of the underlying W, with a confidence level $\gamma \in (0,1)$, is equal to the infimum of x which is a member of the set of real numbers such that the probability of $\mathrm{W} + x$ being less than 0 is equal to or less than $1 - \gamma$.

This paper makes use of a type of historical VaR. As mentioned by Pérignon and Smith (2010), historical VaR is quite intuitive and easy to calculate on a portfolio level since it does not assume a particular parametric distribution of W. According to Pérignon and Smith (2010), historical VaR is the most popular VaR method in the world, as 73% of banks prefer to use historical VaR rather than alternative VaR methods. The downside of historical VaR is that the data needs to be gathered and available. Also, when one uses historical VaR it is important to keep in mind that there exists a risk of underestimating future variance since historical VaR assumes that distributions realized historically also

[19] Stock & Watson (2012) (p. 575).
[20] Jorion (2006) (p. 8).
[21] Jorion (2006) (p. 9).
[22] Jorion (2006) (p. 10).
[23] A similar equation can be found in Artzner et al. (1998) (p. 13.).

will be realized in the future. For example, during times of crises, VaR measures tend to increase. Regardless of the potential downsides, we believe a form of historical VaR is most appropriate for our framework. This is because, as Pérignon and Smith (2010) mention, parametric VaR methods are hard for practitioners to implement in practice. Many banks deal with thousands of risk factors and therefore choose not to attempt to estimate time-varying volatilities and covariances for risk factors. A historical VaR measure is suitable for a practitioner approach. Nordea's Wealth Management also uses a form of historical VaR.

By applying VaR in our creation of optimal portfolios, and not a CRRA utility function, we are relaxing Samuelson's first assumption. Despite the fact that the VaR approach is not explicitly a utility function, and is hence limited with respect to its contribution to the academic debate, it is something we believe is valuable to a market actor.

3.2.1 VaR optimization method

In our application, we optimize portfolios by maximizing expected terminal wealth over different investment horizons given specific constraints. Specifically, we decide the optimal initial allocation for the period from the standpoint of maximal wealth at the end of the period.

We use overlapping periods rather than distinct. For example, for 10-year periods, instead of optimizing over 1981-1990 and 1991 to 2000 we optimize over 1981-1990, 1982-1991, 1983-1992 etc. The main reason for using overlapping rather than distinct periods is to gain more observations. Using 10-year periods we would only have 21 distinct periods compared to 206 overlapping periods. Also, when using distinct periods the choice of initial starting year can have a big impact on the results. The negative aspects of overlapping periods are made up of the correlation between the periods and the fact that it underweights the earliest and latest return periods. The correlation issue will make statistical claims less reliable. With regards to the underweighting, it is best explained with an example; the first and last year in the data set will only be used once in an optimization using investment periods of 10 years, while the middle periods in the same optimization will be used 10 times. This means that the results produced will be slightly biased towards the middle periods. We do not believe the underweighting will have any notable effect on our results, however, we acknowledge the correlation issue and that one should have it in mind when interpreting our results.[24]

Formally, for a given VaR profile our optimization model is defined as follows:

$$\max_{\boldsymbol{w}} E[W(\boldsymbol{w})] = \frac{\sum_{t=1}^{T} W_t(\boldsymbol{w})}{T} \tag{9}$$

$$where\ W_t = \boldsymbol{w} \cdot \boldsymbol{r}_t^T \tag{10}$$

$$Subject\ to \sum_{i=1}^{N} w_i = 1\,, 0 \leq w_i \leq 1, \tag{11}$$

[24] We are aware that the unusually large negative return resulting from the financial crisis of 2008 will be weighted relatively less for longer horizons. On the other hand, we note that 8 out of the 10 periods with the lowest stock returns are counted in full, including those relating to the great depression of 1931.

$$Probability[W_t > (1 - z)] \geq \gamma \quad (12)$$

In equations 9 and 10, W_t is the end of period wealth at time t, $\boldsymbol{w}$ represents a vector of weights (w_i) in the different asset classes while $\boldsymbol{r}_t^T$ denotes the transposed vector of returns for time t. T represent the total number of overlapping rolling periods while N corresponds to the number of asset classes. The γ and z values are the confidence level and maximum loss percentage respectively for the specific VaR case. Equation 11 ensures that the weights of the asset classes must sum to one and that we are not allowed to short any of the assets, and equation 12 represents the VaR restriction.

3.3 Utility functions

Apart from using VaR as a tool when constructing optimal portfolios, we incorporate expected utility functions. Despite the issue of limited practical significance for some cases, i.e. no single utility function is applicable to all agents; expected utility has – ever since Bernoulli (1738) – been fundamental when determining the value of an asset.

As highlighted by scholars such Kritzman and Rich (1998), Stangeland and Turtle (2001) and Kritzman (1992), the allocation of risky assets with respect to investment horizon is sensitive to the utility function. Due to this we optimize using different utility functions. One of these functions serve as our primary utility function, on which we run a number of tests. Other utility functions are used mainly in order to provide perspective to our primary specification.

Our chosen primary utility function is a CRRA utility function, corresponding to the one used by Hanna and Chen (1997). We decided to use a CRRA utility function because we believe, in line with scholars such as Merton (1969) and Janeček (2004), that CRRA is the most general, and hence most reasonable, characteristic of a utility function when it comes to approximating real world agent behaviour.

Formally, our primary utility function is the following:[25]

$$U(W) = \frac{W^{1-\alpha}}{1-\alpha} \; \forall \; \alpha \neq 1, \quad (13)$$

$$U(W) = ln(W) \; \forall \; \alpha = 1. \quad (14)$$

Relative risk aversion (RRA) is defined mathematically as $R(W) = -WU'' \,/\, U'$ which, when applied to equation 13 and 14, gives

$$RRA(W) = \alpha. \quad (15)$$

In equation 13 and 14, W represents total wealth while equation 15 confirms that α represents the level of relative risk aversion.[26] Alto note that when $\alpha = 1$, this utility functions is analogous to the log utility function in table 2. Choosing this framework enables us to test different levels of relative risk aversion. In turn, such a feature allows us to easily describe the effect investment horizon has on equity allocation for a number of different investors characterised by CRRA. According to Janeček (2004), there is a relatively

[25] A formula analogous to the one in equation 9 can be found in Hanna and Chen (1997) (p. 19).
[26] A derivation of equation 15 can be found in equation A1 to A8 in the appendix.

broad consensus around a level of relative risk aversion of between 2 and 10. Cagetti (2003) estimates that the coefficient of risk aversion is usually higher than three and quite often higher than four.

Having introduced our primary CRRA utility function we are looking to compare it to utility functions that are characterised by DRRA and IRRA. We think these tests will be interesting since there is no consensus as to which one of these three measures is the "correct" one. As mentioned, Merton (1969) finds the CRRA to be the most plausible of the three, while Thorley (1995) claims that the observed behaviour of investors together with conventional wisdom suggests that the DRRA best captures preferences over different investment horizons for rational investors. When we asked Nordea, we were told that their clients tend to show a greater aversion to risk as their portfolio increased in value, i.e. that they are characterised by IRRA.

In line with the contents of table 1, one should expect the DRRA function to produce equity allocations that are increasing with investment horizon, relative to CRRA. The opposite should be the case for IRRA: the equity allocation should in this case decrease with investment horizon, relative to CRRA. The utility function we use for this test is an extension of our primary utility function in equation 13, and similar to the utility function used by Thorley (1995).

The utility function is the following:

$$U(W) = \frac{(W-\theta)^{1-\alpha}}{1-\alpha} \; \forall \, \alpha \neq 1. \tag{16}$$

$$U(W) = ln(W-\theta) \; \forall \, \alpha = 1 \tag{17}$$

The function is the same as our primary function with the exception of the utility parameter θ. Applying RRA in the same way as we did to equation 13 and 14 we get:[27]

$$RRA(W) = \frac{\alpha}{1-(\theta/W)}. \tag{18}$$

As we can see from equation 18, θ is a parameter by which we can decide whether the agent has a utility function with CRRA, DRRA or IRRA. If $\theta = 0$ then the utility functions in equation 16 and 17 are analogous to the utility functions in equation 13 and 14, respectively, which means that the agent is characterized by CRRA. If $\theta > 0$, then the investor is characterized by DRRA, and if $\theta < 0$ the investor has an IRRA utility function. It is worth noting that the utility function described by equation 16 and 17 exhibit decreasing absolute risk aversion (DARA) with compatible levels of the parameter θ. This means that when the agent becomes wealthier, he or she will invest a higher absolute amount in risky assets.[28]

When optimizing utility functions with DRRA and IRRA we are relaxing Samuelson's second assumption.

3.3.1 Utility optimization method

In our application, we optimize portfolios by maximizing expected utility, which is a function of wealth, over different investment horizons given specific constraints. Specifically, we decide the optimal initial

[27] A derivation of equation 18 can be found in equation A9 to A18 in the appendix.
[28] A derivation of this result can be found in equations A19 to A22 in the appendix.

allocation for the period from the standpoint of maximal utility at the end of the period. As with the VaR optimization, we optimize using overlapping periods.[29]

For a given investment horizon and level of risk aversion, our optimization model is formally defined as follows:

$$\max_{\boldsymbol{w}} E[U(\boldsymbol{w})] = \frac{\sum_{t=1}^{T} \frac{(W_t - \theta)^{1-\alpha}}{1-\alpha}}{T} \tag{19}$$

$$where\ W_t = \boldsymbol{w} \cdot \boldsymbol{r}_t^T \tag{20}$$

$$Subject\ to \sum_{i=1}^{N} w_t = 1\,, 0 \leq w_i \leq 1 \tag{21}$$

In the above equations W_t is the end of period wealth for time t, θ is the parameter affecting the RRA characteristic, α is the relative risk-aversion coefficient, $\boldsymbol{w}$ represents a vector of weights (w_i) in asset classes while $\boldsymbol{r}_t^T$ denotes the transposed vector of returns for time t. T represent the number of overlapping rolling periods and N corresponds to the number of asset classes. The restricting equation 20 tells us that the weights of the assets all have to sum to one, and that we are not allowed to short any of the assets. The θ parameter takes on zero for the CRRA case, is positive for DRRA and negative for IRRA.

3.4 Optimization and genetic algorithms

In order to find optimal portfolios for our chosen frameworks and time horizons, we make use of optimization programs and algorithms. As for most optimization problems, we work with objective functions, general constraints, and bounds for the variables. Our primary challenge concerns the potential existence of local optima, which typically arise due to nonlinearity in either the objective function or the general constraints. The overall challenge for the optimization procedure is being able to distinguish between a local and a global optimum. Specifically, this is because the program or algorithm can recognise whether the obtained solution is optimal within a neighbouring set of solutions, i.e. a local optimum; but at the same time it is unable to recognise whether this optimum is the optimum of the optimums, i.e. the global optimum. [30] The existence of this challenge will be taken into account in our application.

When tackling optimization, a straight forward method would be to use the Solver function in Excel. The nonlinear optimization program in Solver optimises by using Generalized Reduced Gradient (GRG) code.[31] According to Lee et al. (2004), the GRG method has been proven to be a relatively precise and accurate method for solving nonlinear optimization problems.[32] Like virtually all other optimization methods, the GRG can easily find a local optimum. This means that when Solver communicates to the user that it has found a solution, it means that the GRG code has found a local optimum - but not

[29] The rationale for optimizing over overlapping periods rather than distinct can be found in section 3.2.1.
[30] Chinneck (2006) (ch. 16, p. 1).
[31] Frontline Systems (2017) (a).
[32] The basic concept of the GRG method is thoroughly explained in section 2 in Lee et al. (2004).

necessarily a global optimum.[33] One popular method of mitigating this limitation is to initiate Solver using a number of different starting values. Such an approach would, after a number of tries, ideally make Solver hit all of the available local optima. The user would then choose the local optimum with the most desirable result, which in our case would be the highest expected utility or expected return, which would consequently be considered the global optimum. Apart from the limitation of local optima generally inherent in nonlinear programming, Solver also produces an output with very limited information about the optimization procedure. This is quite natural since Solver is in its core a product designed for an end user with a relatively basic degree of technological sophistication. Such a user will most likely not be negatively affected by the black box characteristic of Solver.

For our application, we prefer not to rely on Solver, specifically due to the issues surrounding local optima, as well as the lack of control of, and insight into, the optimization procedure. Consequently, we have decided on using a programming procedure built on genetic algorithms. A genetic algorithm is an optimization technique inspired by natural selection, and it is used to solve a variety of complex optimization problems, including nonlinear ones.[34] The approach was invented in the 1960s by John Holland, who has since written extensively about the topic.[35]

In deciding on our optimization approach we have taken inspiration from Holland (1992) and Melanie (1998), who both provide various examples of genetic algorithms and their applications.[36] Hence, the following five step model is not taken from a particular source but is a product of our own intuition together with information from the mentioned sources. The genetic algorithm method, applied to our optimization problem, works as follows:

1. Begin by creating one portfolio for each asset class, made up of 100 percent of that specific asset class, i.e. one portfolio each for stocks, Bonds, Bills, gold, and cash. Together these portfolios make up what we call the first generation of portfolios, each portfolio representing an individual.
2. In the next step, pair the different individuals in order to produce an offspring. All possible pairwise combinations will be used, resulting in 10 new individuals (children). The new individuals consist of a random mix from each parent – e.g. 60 percent from one parent and 40 percent from the other.
3. Next, a fitness function is applied, either the utility function or the expected return depending on the problem at hand, to evaluate the five parents plus the 10 offspring. The five best individuals survive and the rest are terminated.
4. Step 2 and 3 are then replicated with the five surviving individuals serving as the new parents, i.e. the second generation of individuals.
5. The process stop when a specified termination requirement is satisfied. Our termination requirement is that the process has to reach 100 generations, after which it stops.

33 Frontline Systems (2017) (b).
34 Melanie (1998) (p. 21).
35 Melanie (1998) (p. 65).
36 Another simpler approach can be found in Melanie (1998) (p. 103).

To improve the model and make it more reliable, a mutation tool is incorporated. Essentially, a mutation tool serves to prevent the model from getting stuck on a local optima.[37] In or model, mutation plays a role in the second step when the pairing is done. Instead of simply pairing two parents, random mutation occurs. This means that for each asset class, there is a small random change (we have applied both a percentage mutation as well as a percentage point mutation). We terminate our model when the results clearly converge and does not change. We find that this happens after 50-100 generations.

The genetic algorithm serves as an intuitive instrument for creating optimal portfolios. In essence, it starts off with the assumption that all asset classes are equally likely to be relevant, then continuously creates new portfolios in order to see if a better alternative can be found. The process is based on randomized numbers to create random combinations of the different portfolios, which is why a given procedure might produce different results, depending on the randomization. However, given an infinite amount of generations, the result from the process should in its limit reach the correct optimal portfolio. While an infinite number of generations is not possible in a practical sense, it is that exact idea of the limit that makes us use as many as 100 generations, even though the "best" portfolio might already be identified, which it in many cases is, after around 50 generations.

Lastly, we find that the results produced by our genetic algorithm match the solutions provided by Excel's solver for the first and often second decimal for the percentage allocation to stocks. The findings and conclusions we report in this paper are not sensitive to the second or even the first decimal of the percentage allocation, and hence Excel's solver and our genetic algorithm work to confirm each other.

---

[37] Melanie (1998) (p. 129).

4. Data

This paper uses a data set consisting of historical market returns originally obtained by Jeremy Siegel (2016).[38] Returns for the most previous years, 2014-2016, have been retrieved from Thomson's Datastream. The complete data set is made up of U.S. market data for the years 1802-2016, which makes it one of the most extensive sets of market data with respect to the U.S. The data set consists of nominal return data for stocks, Treasury bonds, Treasury bills and gold. The stock returns cover publicly listed firms and take survivorship bias into account, i.e. the returns also cover firms that have disappeared during the time period. Siegel (2014) points out that certain early historical values are unavailable and that these have been adjusted for accordingly. For example, Treasury bonds have been unavailable for some of the early periods, in those cases high-grade municipality bonds, for which default premiums have been estimated and removed, have been used as proxies. The change in the consumer price index (CPI) is also included for the entire period. Using the CPI, we converted the nominal returns into real returns and constructed return data for cash. The average annual inflation in our data for the full period has been 1.56 percent. A summary of the real return data can be found in table 3.

Table 3 – Summary statistics of real returns, 1802-2016

	Stocks	Bills	Bonds	Gold	Cash
Mean	8.31%	2.82%	3.90%	1.27%	-1.23%
Std.dev	18.08%	5.92%	8.77%	13.29%	5.45%
Median	7.69%	2.54%	3.88%	0%	-1.21%
Maximum	66.62%	23.68%	35.13%	99.94%	17.98%
Minimum	-38.57%	-15.63%	-21.86%	-38.13%	-21.28%
Skewness	0.0849	0.2122	0.2689	2.8304	-0.1295
Kurtosis	0.6201	2.1509	1.0554	17.473	2.4378

The Siegel data set was made available to us by Nordea together with the expressed wish that we should use it. However, even if the data set would not have been supplied to us we would still have opted for a similar, if not the same, data set. This is mainly due to the fact that the data set includes a long time horizon which is extra valuable for our application. Essentially, time diversification analysis quite heavily relies on the potential existence of long term trends. Having a data set that includes most of financial history also enables us to detect possible paradigm shifts in the return structure of our chosen assets. The set is also preferable to us since U.S. market data, echoed by Bianchi et al. (2016), serves as a good proxy for aggregate international data. Bianchi et al. (2016) get similar results from using exclusively U.S. data as they do when using data from a 20-country average. We believe the positive aspects gained from a high degree of generality outweighs the negative aspects of not completely corresponding to an individual country such as Denmark.

We find the data set to be both valid and credible. Partly because the data set is recognised and used by scholars, but also because Jeremy Siegel is professor at Massachusetts Institute of Technology (MIT) and a

[38] Siegel (2016).

widely known financial economist. According to Siegel (2014) the data set has been constructed by merging data from different sources. The sources include Goetzmann-Ibbotson, the Cowles Foundation and the Center for Research in Security Prices. We find all of these sources to be credible on their own. Concerning our addition of data for 2014-2016, Thomson's Datastream is a large acknowledged financial statistical database, and hence we consider it a reliable source of data.[39]

[39] Access to Thomson Datastream was provided by the library at Copenhagen Business School.

5. Results

5.1 Returns and risk measures over time and for different investment horizons

We begin by graphing the historical real returns of asset classes in our data set. Observing figure 1 we can see that over the last two centuries the total real return on stocks dominate all other asset classes. If an agent invested $1 in stocks in the beginning of 1802, and continuously reinvested the capital gains and dividends, the $1 would have grown to $1,258,548 in real terms in 2016. If the same agent would, in the beginning of 1802, have invested his or her $1 in bonds, bills, gold or cash the $1 would in 2016 have had a real value of $1,756, $274, $3 and 0.05$ respectively.

Figure 1 – Historical real returns for different asset classes

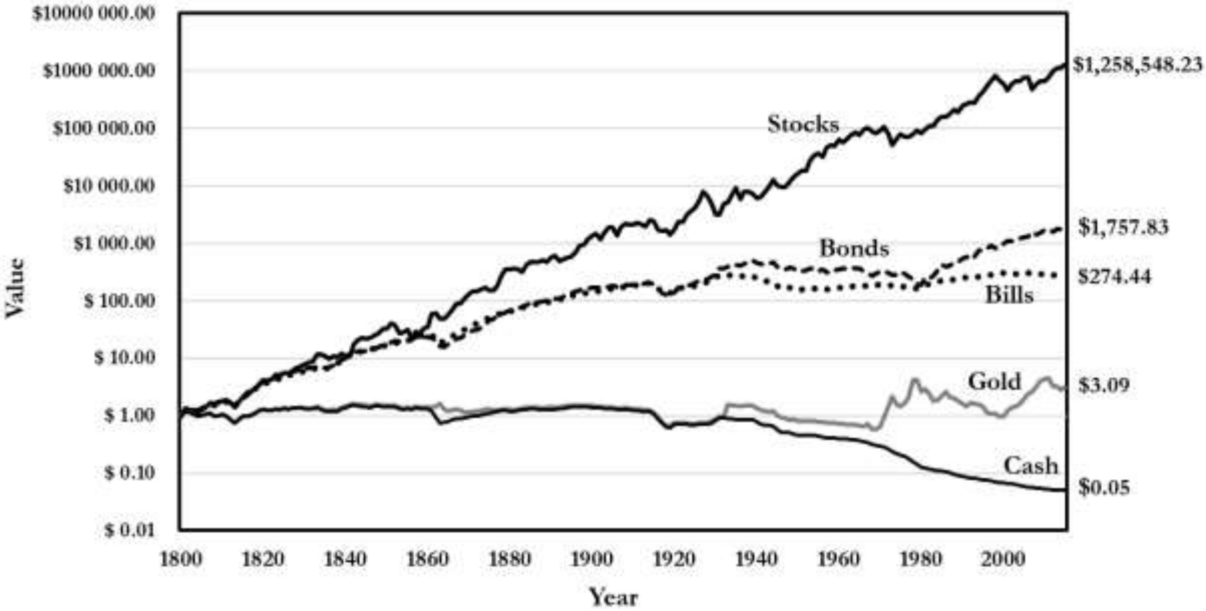

The figure illustrates characteristics of U.S. data for the years 1802-2016. Logarithmic scale. Initial wealth is $1. Real returns are depicted.

One reason for the superior historical performance of stocks in figure 1 is inflation risk. During the inflationary decades of the 1940s, 1950s, 1960s and 1970s, bonds had a negative real rate of return while stocks had positive real returns[40]. The restoration of price stability during the last decades has led to positive real return on bonds, and to a lesser extent bills.[41] This is in line with Bodie (1995) and Siegel (2014) who highlight the exposure of bonds to inflation risk as well as Cagan (1974) who finds that stocks have performed well relative to inflation.

Figure 1 give us an early indication that time diversification is prevalent and that an investor should invest all of his or her financial assets in equities if the investment horizon is sufficiently long. From an initial glance, the two time diversification tenets by Dorsett and Reichstein (1995) seem to be valid. However, Samuelson (1992), being aware of historical return data for stocks and other asset classes, points out that we only have one history of capitalism. He points out that inferences based on one sample must never be taken as a certainty, regardless of the size of the sample. However, despite Samuelson's critique, scholars as well as market practitioners believe figure 1 provide a strong argument for time diversification.

[40] Sowell (2014) (p. 567).
[41] Sowell (2014) (p. 568).

As a next step, after having observed the historical real returns, we would like to see if risk for different investment horizons also give an indication of time diversification.

Figure 2 – Interval of real returns for different investment horizons.

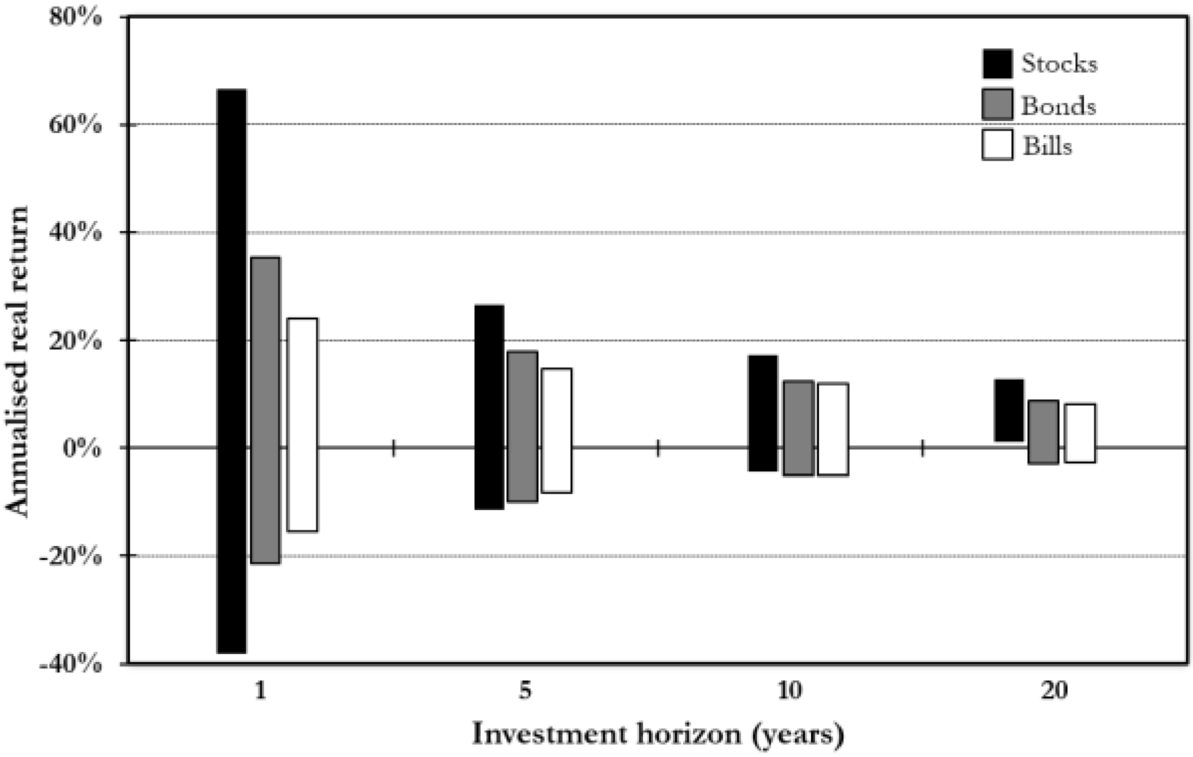

The figure illustrates the range of return outcomes for stocks, bonds and bills using U.S. data for the years 1802-2016. Results are shown for investment horizons of 1, 5, 10 and 20 years.

Figure 2 depicts the range of the annualized real return for each asset based on holding period. For an investment horizon of one year, we can see that the highest and lowest real return on stocks are significantly further apart than the highest and lowest real return for bonds and bills. However, when we increase the investment horizon to five years the worst performance of stocks is almost the same as the worst performance of bonds and bills. Lengthening the investment horizon to 10 years, we can see that the worst performance of stocks is better than that of bonds and bills. For 20 years, the trend is even more dominant while all possible outcomes for stocks are above zero.

Figure 2 stands in contrast to the notion that it is riskier for an investor to invest all of his or her financial wealth in stocks rather than bonds or bills for a long investment horizon. These results are in line with Lebowitz and Kogelman (1991) who find that the volatility inherent in equities are mitigated for long horizons. If an investor were looking to save for retirement while assuming that future returns of assets will look somewhat similar to the past, then it would, from looking at figure 2, seem rational for the investor to invest most or all of his or her financial assets in stocks. This strengthens the notion that the two time diversification tenets by Dorsett and Reichstein (1995) hold true.

Having observed figure 2, we continue by taking a closer look at different risk measures for stocks and how these are affected by the investment horizon.

Table 4 – Risk measures of total stock returns for different investment horizons

Years	1	2	3	4	5	10	15	20
Mean	0.0831	0.1717	0.2625	0.3607	0.4647	1.0758	1.9861	3.2102
Std.dev.	0.1809	0.2766	0.3461	0.4216	0.4917	0.8411	1.4282	2.0875
Variance	0.0327	0.0765	0.1198	0.1778	0.2418	0.7074	2.0396	4.3577
Range	1.0519	1.5189	1.9859	2.5430	2.7052	4.0900	6.5897	9.5896
Best return	0.6662	0.9872	1.3852	1.9470	2.2634	3.7475	6.3580	9.8195
Worst return	-0.3857	-0.5317	-0.6007	-0.5960	-0.4418	-0.3424	-0.2317	0.2299
Semideviation	0.1799	0.2551	0.3063	0.3697	0.4344	0.7100	1.1783	1.5679

Table 4 includes mean, standard deviation, variance, range, best return, worst return, and semideviation for stock returns at different investment horizons. The mean return is included, even though it is not a risk measure in itself, because it serves as the primary reason for investing. Unsurprisingly, the mean return increases as the time horizon increases. The same is the case for the standard deviation and variance. The range of the best and worst return also increases as the time horizon grows. Semideviation, which is the standard deviation of all the observations that are below the mean, is not substantially different from the standard deviation for all observations. This tells us that the variation for outcomes below the mean are not much different from outcomes above the mean.

Looking at the risk measures of total stock returns it is evident that these give a different picture than the one we get from looking at figure 1 and figure 2. However, the results from table 4 provide a misleading picture given that they express measures for the total returns rather than the annualized returns. Hence, we go on to look at the risk measures once again, however this time we look at the annualized returns.

Table 5 – Risk measures of annualized stock returns for different investment horizons

Years	1	2	3	4	5	10	15	20
Mean	0.0831	0.0748	0.0716	0.0703	0.0695	0.0676	0.0681	0.0686
Std.deviation	0.1809	0.1290	0.1000	0.0841	0.0732	0.0449	0.0346	0.0260
Variance	0.0327	0.0166	0.0100	0.0071	0.0054	0.0020	0.0012	0.0007
Range	1.0519	0.7254	0.5997	0.5130	0.3769	0.2096	0.1597	0.1160
Best return	0.6662	0.4097	0.3361	0.3102	0.2669	0.1685	0.1423	0.1264
Worst return	-0.3857	-0.3157	-0.2636	-0.2028	-0.1101	-0.0411	-0.0174	0.0104
Semideviation	0.1799	0.1286	0.1010	0.0863	0.0762	0.0468	0.0366	0.0275

Table 5 comprises the same risk measures as table 4, with the exception that the measures are for annualized returns instead of total returns. The mean annualized stock return decreases slightly with time, due to the use of geometric averages. The standard deviation and subsequently the variance is more clearly affected as they decrease substantially with investment horizon. Additionally, the range between the best and the worst outcome decreases substantially with time, while the worst outcome at the same time improves with time to the point that it is in the positive, as also depicted in figure 2. As with the standard deviation, the semideviation decreases with the investment horizon.

The risk measures of annualized stock returns differ from the risk measures of total returns in the sense that they decrease significantly with the length of the investment horizon. Similarly to figure 1 and figure 2, this table support the notion of time diversification and is in line with the tenets by Dorsett and Reichstein (1995).

5.2 Autocorrelation in stock returns

Given the results in section 5.1, particularly the ones in figure 1, a natural next step is to test for mean reversion in stock returns. In order to test for mean reversion, we run an autoregressive (AR) model.[42] Initially we test with lags between 1 and 10 in order to check for notable patterns. When conducting these tests, we find that lag two and lag five are consistently significant on the five percent level while being negative. In all tests, lag 2 has the higher significance of the two, with a p-value lying steadily around 0.01. Lag 5 has a p-value slightly below 0.05.

After reviewing the different AR specifications, we chose the model with five lags, i.e. the AR(5). We chose the AR(5) because it represented the best trade-off between the result from the Akaike information criteria (AIC) and our own intuition. Would we have relied fully on the recommendation given by the AIC we would have ended up with an AR model with two lags, while the second best model according to AIC is the AR(5). Opting for an AR model with two lags does not seem realistic since it would imply that the stock market recovers (relapse) within two years after a market downturn (upturn), whereas the AR(5) model allows for five years.[43] In their analysis, Blanchett et al. (2013) also choose an AR(5) model. However, they admit that their selection of an AR(5) model is somewhat subjective.

Table 6 – SAS results from AR(5) model

Parameter Estimates				
Variable	Estimate	Standard Error	t Value	Pr > |t|
Intercept	0.1197	0.0189	6.32	<0.0001
Lag1	-0.0080	0.0693	-0.12	0.9082
Lag2	-0.1850	0.0691	-2.68	0.0081
Lag3	-0.0616	0.0702	-0.88	0.3807
Lag4	-0.0454	0.0692	-0.66	0.5125
Lag5	-0.1396	0.0690	-2.02	0.0444
Ordinary Least Squares Estimates				
MSE	0.0319		DFE	204
MAE	0.1364		Root MSE	0.1787
MAPE	202.8130		AIC	-121.3430
ESS	0.3560		TSS	6.879
SSE	6.5160		R-Square	0.0518

Table 6 depict the results of an AR(5) model in SAS. We can see that all lags have a negative coefficient, and as mentioned previously, the second and fifth lag are significant on the five percent level. Negative

[42] The AR model, AIC and our application is described in section 3.1.
[43] The SAS results for the AR(2) model can be found in table A1 in the appendix.

lags indicate mean reversion, and the fact that all lags, particularly the statistically significant ones, are negative, supports the notion of time diversification. To put it another way, extreme events smooth out over time because negative coefficients resulting in low returns tend to be followed by high returns and vice versa.

The existence of mean reversion in stock returns goes against Samuelson's second assumption, i.e. that stock returns are independently and identically distributed (IID). Since the second assumption is violated, is it then safe to claim that an investor should increase his or her allocation to equities as the investment horizon increases? According to Stangeland and Turtle (1999), if the return process of the asset exhibits mean reversion, a constant relative risk aversion (CRRA) investor should increase his or her exposure to risky assets with time. However, looking at table 2 by Kritzman and Rich (1998), an investor with log utility, i.e. with CRRA, should hold the exposure to risky assets constant with time. Additionally, we would like to know if the potential allocation recommendations from this return process could be compared to the VaR results obtained by Panyagometh (2011). We therefore go on to create optimal portfolios by using the historical data in combination with our value at risk (VaR) and expected utility frameworks.

5.3 Optimal equity allocations using the VaR framework

The first of our two primary specifications is the VaR specification.[44] We chose to optimize our framework using different VaR levels, which are the same as the levels used by Nordea's Wealth Management. The returns used in this analysis are nominal returns due to nature of the VaR. If we were to use real return data the pedagogical feature of the VaR approach would be mitigated. An investor would have to take inflation into account when deciding on the VaR limit. Furthermore, the standard approach is to use nominal returns when using VaR, for example, Nordea uses nominal returns. The VaR levels correspond to different investment profiles: -0 percent (conservative), -7.5 percent (moderate), -15 percent (balanced), -20 percent (growth) and -30 percent (return focus).

[44] A description of our specification and the justification we have for using VaR as a risk measure can be found in section 3.2.

Figure 3 – Optimal stock allocations for different VaR-levels, 95 percent certainty

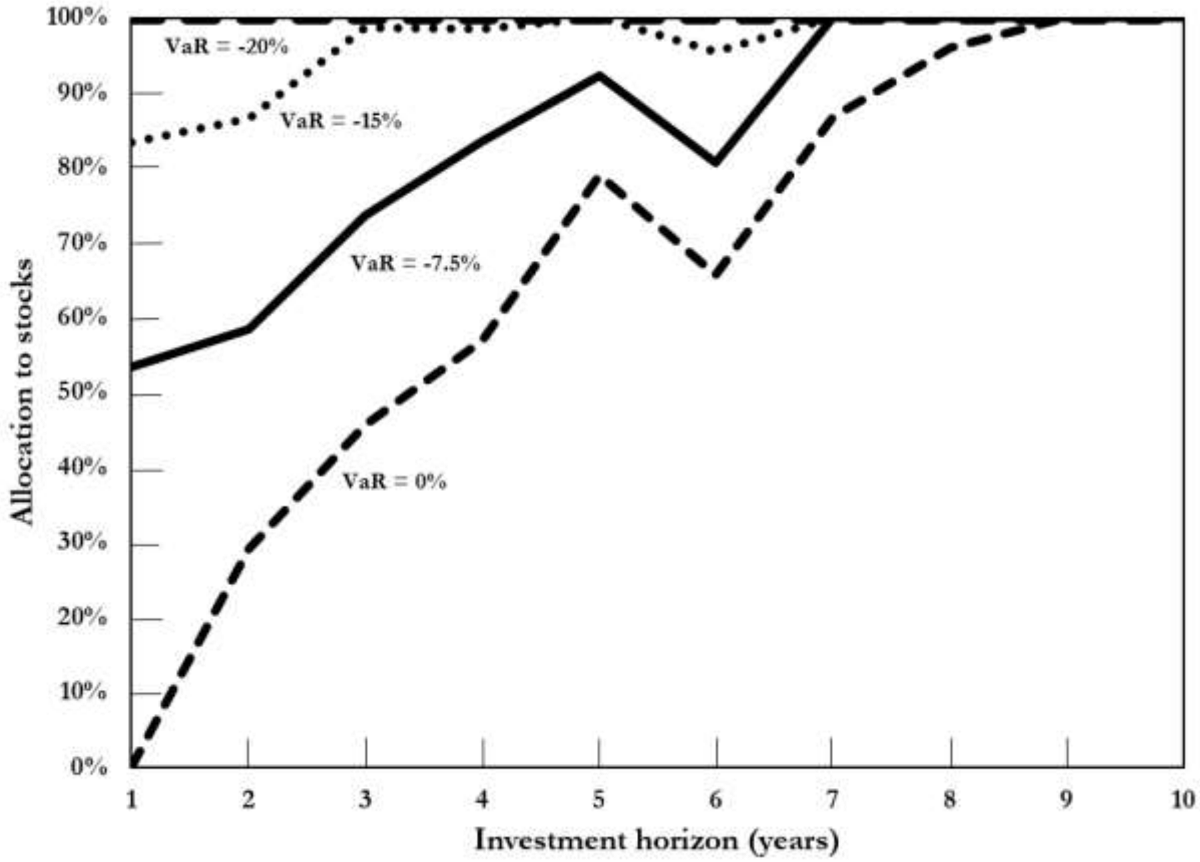

Framework is optimized over investment horizons ranging from 1 to 10 years, with one-year increments. Five different VaR-levels are used, -30 percent, -20 percent, -15 percent, -7.5 percent and 0 percent. Each graph depicts the optimal allocation for a specific VaR-level. Nominal returns are used.

Figure 3 shows the optimal allocation to equities given different investment horizons and VaR-levels for a 95 percent level of certainty.[45] Looking at the two least conservative investment profiles, the -20 percent (growth) and -30 percent (return focus), we can see that our specification recommends a 100 percent allocation to equities regardless of the length of the investment horizon.[46] This makes us unable to see if there is a dime diversification aspect for these allocations, since equities are superior to all other assets from year one and onwards. For the three more conservative profiles, -0 percent (conservative), -7.5 percent (moderate) and -15 percent (balanced), the specification recommends an increasing allocation to equites as the investment horizon increases.

One factor that stands out is the downward kinks at the six-year horizon, which we do not expect to find assuming time diversification exists. We believe this can be caused by a few single years in the data set. With a 95 percent certainty level as in this case, a drop in the return during just a few years can have a significant adverse effect on the equity recommendation. Section 5.9 is devoted to look more into kinks in the output(s).

Having looked at the output for a 95 percent certainty, we would like to know more about the equity allocation for a 99 percent certainty. With respect to the profiles for -20 percent (growth) and -30 percent

[45] The optimization model used for the VaR profiles can be found in section 3.2.1. The full composition of the optimal portfolios can found in tables A3 to A7 in the appendix.
[46] The line for -30 percent (return focus) is excluded since it overlaps perfectly with the line for -20 percent (growth).

(return focus), we would like to see if these have a different intercept and if the recommended allocation will in such a case increase with the investment horizon.

Figure 4 – Optimal stock allocations for different VaR-levels, 99 percent certainty

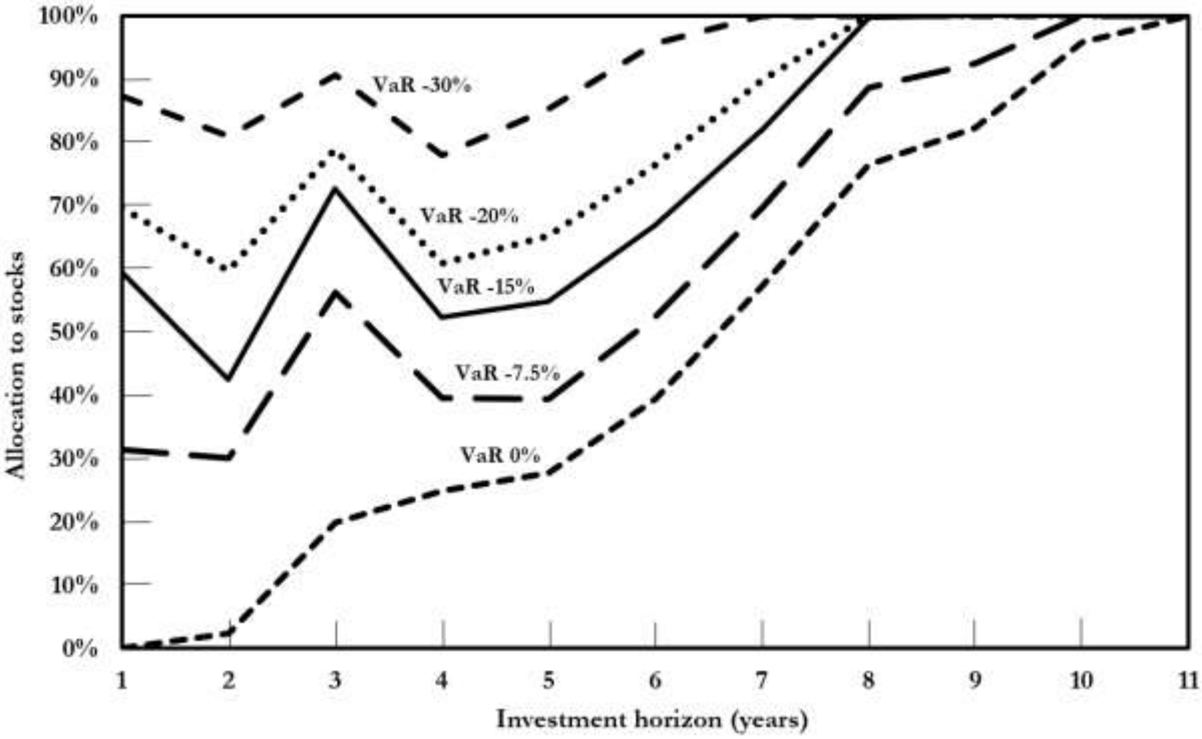

Framework is optimized over investment horizons ranging from 1 to 10 years, with one-year increments. Five different VaR-levels are used; -30 percent, -20 percent, -15 percent, -7.5 percent and 0 percent. Each graph depicts the optimal allocation for a specific VaR-level. Nominal returns are used.

Observing figure 4, which shows the equity allocation for the same five investment profiles as the ones in figure 3, but with a 99 certainty, we can see that all the profiles have an optimal equity allocation of less than 100 percent for an investment period of one year. All the profiles also show an increased allocation to equity with investment horizon, which is to be expected after seeing the results in figure 3. For an investment horizon of 11 years, the framework recommends an equity allocation of 100 percent regardless of investment profile.

Again, we see kinks in the graphs. Instead of the downward kink at an investment horizon of six years in figure 3, the allocations in figure 4 display an upward kink for an investment horizon of three years. As mentioned, we will look more closely into these kinks later in section 5.9.

Having reviewed the optimal equity allocation for frameworks with a 95 percent and 99 percent certainty we would like to see the results of regressing optimal equity allocation on investment horizon.

Table 7 – Linear regressions for different levels of VaR

	99% confidence level	
VaR	Adj. intercept	Adj. slope
0%	-0.0639[47]	0.1085
-7.50%	0.2405	0.0799
-15%	0.4566	0.0610
-20%	0.6004	0.0429
-30%	0.8144	0.0206
	95% confidence level	
VaR	Adj. intercept	Adj. slope
0%	0.1607	0.1151
-7.50%	0.5570	0.0724
-15%	0.8733	0.0245
-20%	1	0
-30%	1	0
	90% confidence level	
VaR	Adj. intercept	Adj. slope
0%	0.2494	0.1599
-7.50%	0.7760	0.0671
-15%	1	0
-20%	1	0
-30%	1	0

Table 6 contains the adjusted intercept and the adjusted slope that best describe the equity allocation regressed on investment horizon. The adjustment indicates that the regressions are not created from using all observations, only the observations up until when the recommendation reaches 100 percent equities. Furthermore, adjusted intercepts refer to a one-year investment horizon instead of zero. Put differently, the intercept refers to the recommended allocation for a one-year horizon, and the slope refers to the suggested increase in equity allocation given an incremental increase of one year to the investment horizon. For each of the three confidence levels, the slope is positive regardless of VaR level. For the less conservative VaR levels for the 95 percent and 90 percent confidence levels, the recommendation is to have a 100 percent allocation to equities regardless of investment horizon.[48]

It is also worth pointing out that the adjusted intercept does not exactly equal the one-year allocation depicted in figure 3 and figure 4. This is because the linear regressions are fitted approximations, i.e. providing a rule of thumb rather than an exact recommendation. When observing the adjusted regressions, it is important to note that these serve as rules of thumb rather than specific recommendations. This is because the inputs in the regressions should be uncorrelated, and this is not the

[47] Given that we apply a framework where shorting is not allowed, a negative value refers to a zero percent allocation to stocks.
[48] The graphical output for 90 percent certainty can be found in figure A1 in the appendix.

case in this framework since the returns used to construct equity allocations for different investment horizons are the same. Also, the use of overlapping periods to obtain equity allocation adds further correlation.

The increasing optimal equity allocation with investment horizon in figure 3, figure 4 and table 6 give us reason to believe that time diversification is prevalent, and that a market practitioner should recommend a higher equity allocation for long term investors than for short term investors. These results are similar to those by Panyagometh (2011), who used a VaR approach on data from Thailand.

However, as Thorley (1995) points out, the chance of stocks providing a negative return decreases with investment horizon since the expected return of equities are positive. Thorley (1995) shows that this is true even when stock returns are IID, i.e. does not exert mean reversion. As mentioned previously in this paper, the VaR approach is similar to a discontinuous utility function, i.e. in the sense that the agent prioritizes not violating the specified VaR level above everything else. Given that the probability of obtaining a negative return from equities decreases with investment horizon, the presented VaR specification is biased towards equities. In general as well as compared to the expected utility framework. In other words, we would expect the resulting graphs from the VaR analysis to be upward sloping even without mean reversion in the data, although mean reversion should lead to more aggressive equity allocation, ceteris paribus.

In order for us to be able to provide more reliable results surrounding the existence of time diversification, we move on to the expected utility framework.

5.4 Optimal equity allocations using the expected utility framework

Figure 5 – Optimal stock allocation for different levels of risk aversion, CRRA utility

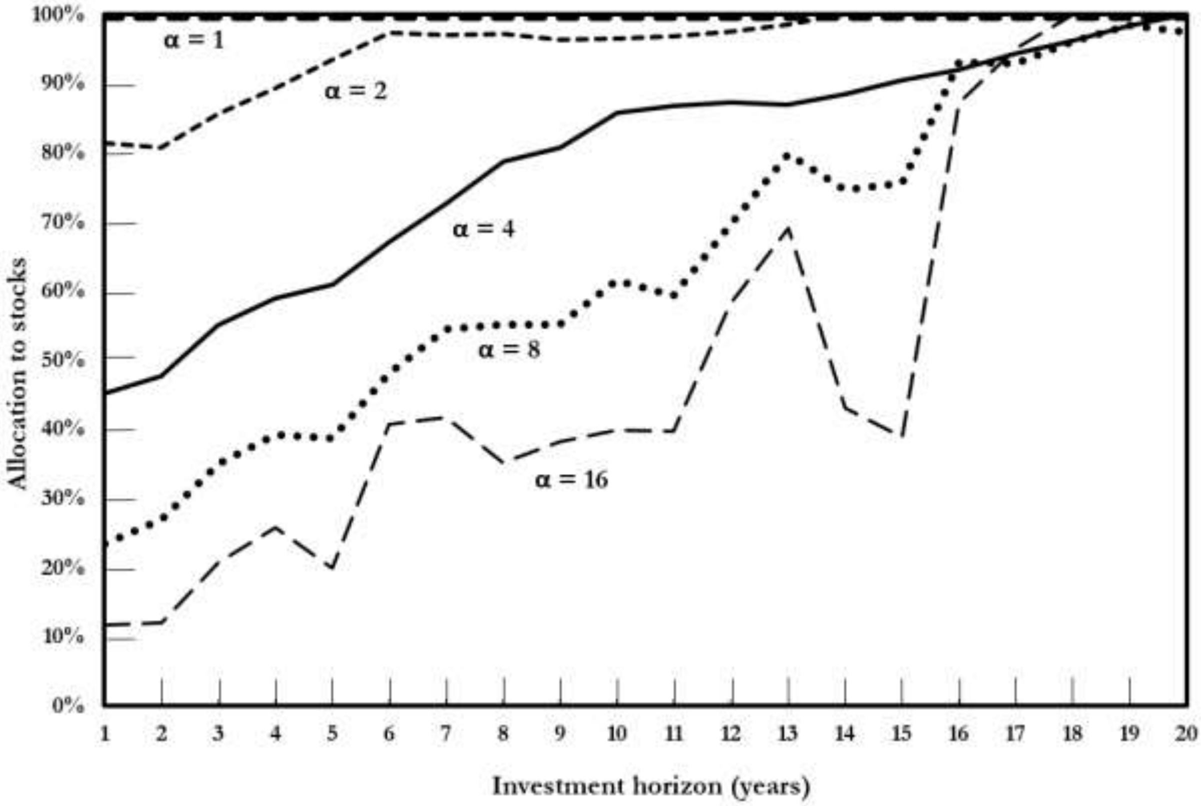

Framework is optimized over investment horizons ranging from 1 to 20 years, with one-year increments. Five different risk coefficients are used; 1, 2, 4, 8, and 16. Each graph depict the optimal allocation for a specific level of risk aversion. Real returns are used.

Figure 5 depicts the optimal equity allocation for different investment horizons for an investor characterized by a utility function with constant relative risk aversion (CRRA).[49] Each of the graphs correspond to a different risk aversion coefficient (alpha), the higher the risk aversion coefficient the more risk averse is the investor. We chose to go with risk aversion coefficients of 1, 2, 4, 8, and 16 since these are in line with the coefficients used by Blanchett et al. (2013) as well as consistent with the coefficients estimated by Janeček (2004) and Cagetti (2003).

Looking at the figure we can see that the recommendation for an investor with an alpha of one is to allocate 100 percent of his or her portfolio to equities for all investment horizons. This result is in line with the predictions in table 2 by Kritzman and Rich (1998), that a log wealth investor should have a constant percentage allocation to equities irrespective of the investment horizon and the return process of the risky asset. Although, since the recommended allocation for one year is 100 percent equities we are unable to say if the allocation would have increased with investment horizon if the one-year allocation would have been less than 100 percent. Looking at the other risk profiles, it would for example be surprising if the recommendation for the log utility investor would been a constant 80 percent allocation to equites.

[49] The utility function that is used is described by equation 13 and 14, while the optimization model for the expected utility framework is found in section 3.3.1. The full composition of the optimal portfolios can found in tables A8 to A12 in the appendix.

The graphs, with an alpha of 2, 4, 8 and 16, which depict the four other risk profiles all show an increasing allocation to equites as the investment horizon increases. Unsurprisingly, the more risk averse agents are recommended to allocate a relatively smaller part of their portfolio to equites. However, even for the most risk averse investor, with an alpha of 16, is recommended to allocate a certain part of his or her portfolio to equities regardless of investment horizon.

As with our VaR framework, the graphs in figure 5 exhibit some kinks. Especially the one for an investor with an alpha of 16, which displays a relatively large downward kink at the 14th and 15th year horizons. It is worth adding that the profile with an alpha of 16 describe a highly risk averse investor, who is therefore very sensitive to years with low returns. Once again, we refer to section 5.9 for a closer look at the kinks.

We would like to look more closely into this framework, specifically to see if there is a positive slope when one regresses equity allocation on investment horizon. When continuing with this specification, both in terms of tests and extensions, we use a risk aversion coefficient of four. There is no consensus as to which coefficient is most credible. Hence, we chose a coefficient of four partly as a result of our own intuition, but also because it lies in line with descriptions by Janeček (2004), as well as estimations by Cagetti (2003). As a practical example, a CRRA individual with a risk aversion coefficient of four is indifferent between an investment opportunity where there is an equal probability between a 15 percent gain and a 9.4 percent loss. In contrast, with a coefficient of 16, only a 4.1 percent loss is required to offset the 15 percent gain. The contrast is even more apparent, when considering a 30 percent gain. In that case, the required loss is 13.5 percent for a coefficient of 4, however, with a coefficient of 16 the loss is 4.5 percent, which is only marginally lower than the previous case.[50]

[50] The numerical examples are derived by fixing the percentage gain and then solving for the percentage loss. The percentage loss must result in the expected utility of the two outcomes equaling the utility from the initial level of wealth. Using a CRRA utility function makes the initial level of wealth irrelevant. In the examples provided, a zero percent return is the certainty equivalent of the investment opportunity.

Figure 6 – Regression of investment horizon on stock allocation

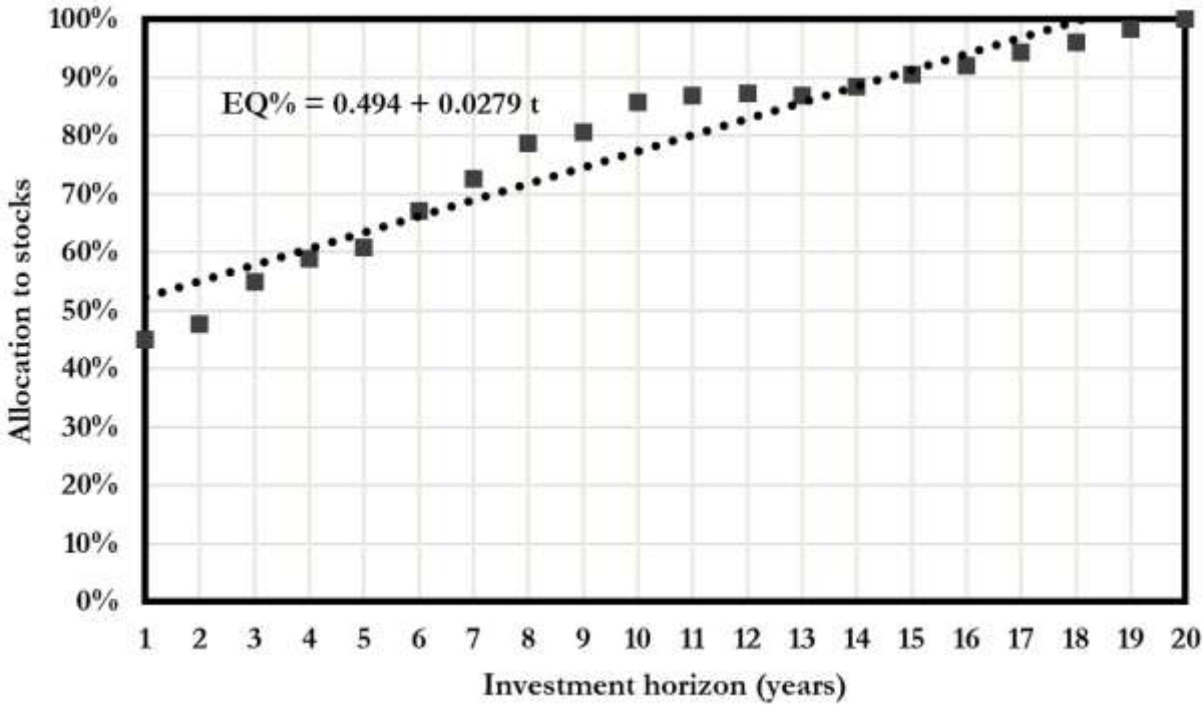

Linear regression depicting stock allocation on investment horizon. CRRA utility with a risk aversion coefficient of four. Each of the larger points represent the optimal stock allocation for a certain investment horizon.

Figure 6 depicts the regression of optimal allocation to equities on investment horizon for an investor characterized by a utility function with CRRA and a risk aversion coefficient of four. According to the regression, the initial allocation to stocks for one year is 52.19 percent, and for each additional year added to the investment horizon, the allocation to stocks should increase by 2.79 percent.

The results from figure 5 and figure 6 provide a strong argument for historical time diversification. Even though Samuelson's second assumption of IDD is violated, due to mean reversion in historical equity returns, both Samuelson's other assumptions are adhered to in this framework. This fact that the investor has CRRA utility and that returns are only a function financial assets, are two things that add weight to the academic validity of these equity recommendation. The results from the CRRA framework is in line with those by Hanna and Chen (1997), Blanchett et al. (2013) as well as Hansson and Persson (2000). Although, the recommended allocation to risky assets is considerably less aggressive than the ones proposed by Dolvin et al. (2010) and Barberis (2000).

Having looked at the results for an investor characterized by CRRA, we would like to test what the optimal equity allocation would be for an investor with a utility function that has decreasing relative risk aversion (DRRA) and increasing relative risk aversion (IRRA). As mentioned in section 3.3, even though most scholars consider CRRA to be the most plausible risk characteristic, Thorley (1995) argues that DRRA best describes an investor, while Nordea's Wealth Management has told us that their investors exhibit IRRA. It is also interesting to see if the expected utility framework is robust to a relaxation of Samuelson's first assumption.

Figure 7 – Optimal stock allocation for DRRA, CRRA and IRRA utility

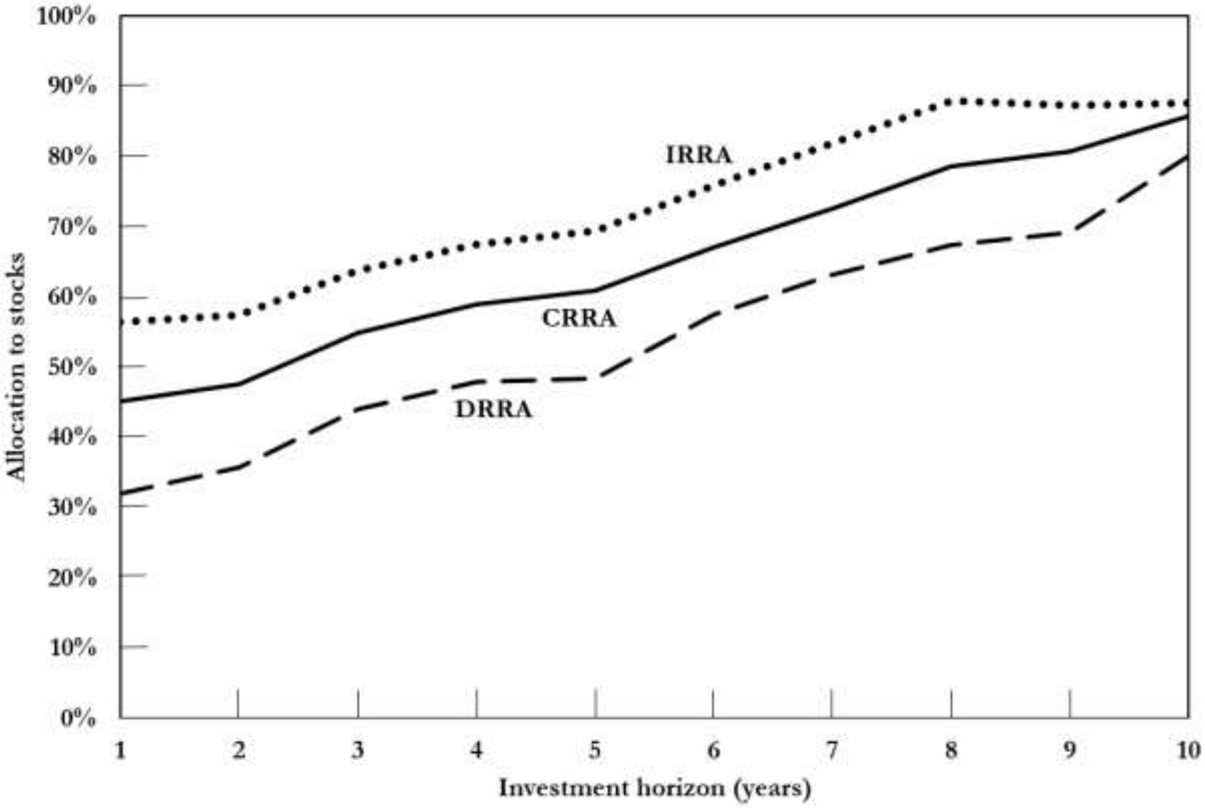

Framework is optimized over investment horizons ranging from 1 to 10 years, with one-year increments. The graphs depict stock allocations for DRRA, CRRA and IRRA utility with a risk aversion coefficient of four. The utility parameters equal to -0.3, 0 and 0.3. Initial wealth is assumed to be $1. Real returns are used.

Figure 7 shows the optimal allocation for an investor characterized by a utility function with either DRRA, CRRA or IRRA, with a risk aversion coefficient of four.[51] The utility parameter, theta, equals -0.3, 0, 0.3. We chose this level on the parameter because it lies in line with scholars such as Thorley (1995), but also because it is compatible with our formal framework.[52] It is worth adding that Thorley (1995) applied the utility parameter to simulated data, and that no other scholar that we know of have optimized utility functions with DRRA and IRRA based on historical data. We note that when applying utility functions with IRRA and DRRA it is important to talk about the initial wealth value. By definition, utility functions exhibiting IRRA and DRRA will have changing preferences, in percentage terms, given different amounts of wealth. Our framework assumed an initial wealth of $1, which makes the application harder to apply directly to a real-life situation. However, this extension is primarily applied to test the robustness of the CRRA application. In addition, there is no single level of starting wealth that is applicable to the typical agent. It could be argued an initial wealth level of $1 is as realistic as any other level.

Looking at figure 7, we can see that the intercept for the DRRA is below the intercept of the CRRA and IRRA. The intercept for the IRRA is above the other two intercepts. All three graphs recommend an increasing allocation to equities with investment horizon. However, the optimal equity allocation increases faster for a DRRA investor than for the other two, and the allocation for a CRRA investor increases faster

[51] The utility function used is described by equation 16 and 17, while the optimization model for the expected utility framework is to be found in section 3.3.1.

[52] Thorley (1995) apply a parameter of 0.7 which isn't compatible with our framework. Looking at the utility function (see section 3.3) it is evident that the parameter must lie below the worst possible outcome of returns for the utility function to make sense.

than the allocation for an IRRA investor. In relative terms, this is in line with the information provided by Kritzman (1992) in table 1. However, even the equity allocation for an IRRA investor increases with investment horizon. Had the return process of stocks been IID, we would expect such an investor to allocate a smaller part or his or her portfolio to equities as the investment horizon increases. We believe that the significant mean reversion characteristic inherent in the data leads to increasing equity allocations, even in the case of IRRA.

By extending the framework to include utility functions characterised by DRRA and IRRA, the application relaxes both Samuelson's first and second assumption. The fact the we relax the second assumption does not provide a major change in the results since the mean reversion characteristic inherent in the historical data provides most of the basis for time diversification. The framework presented in figure 5 thus appears robust to a relaxation of Samuelson's first assumption, giving us further indication of historical time diversification.

Having tested our primary utility framework with respect to Samuelson's first assumption, we would like to test what happens when one relaxes Samuelson's third assumption, the assumption that total wealth is only a function of financial assets.

5.5 Optimal equity allocation with inclusion of non-financial assets

Figure 8 – Optimal stock allocation for different investment horizons and percentages of financial wealth as part of total wealth

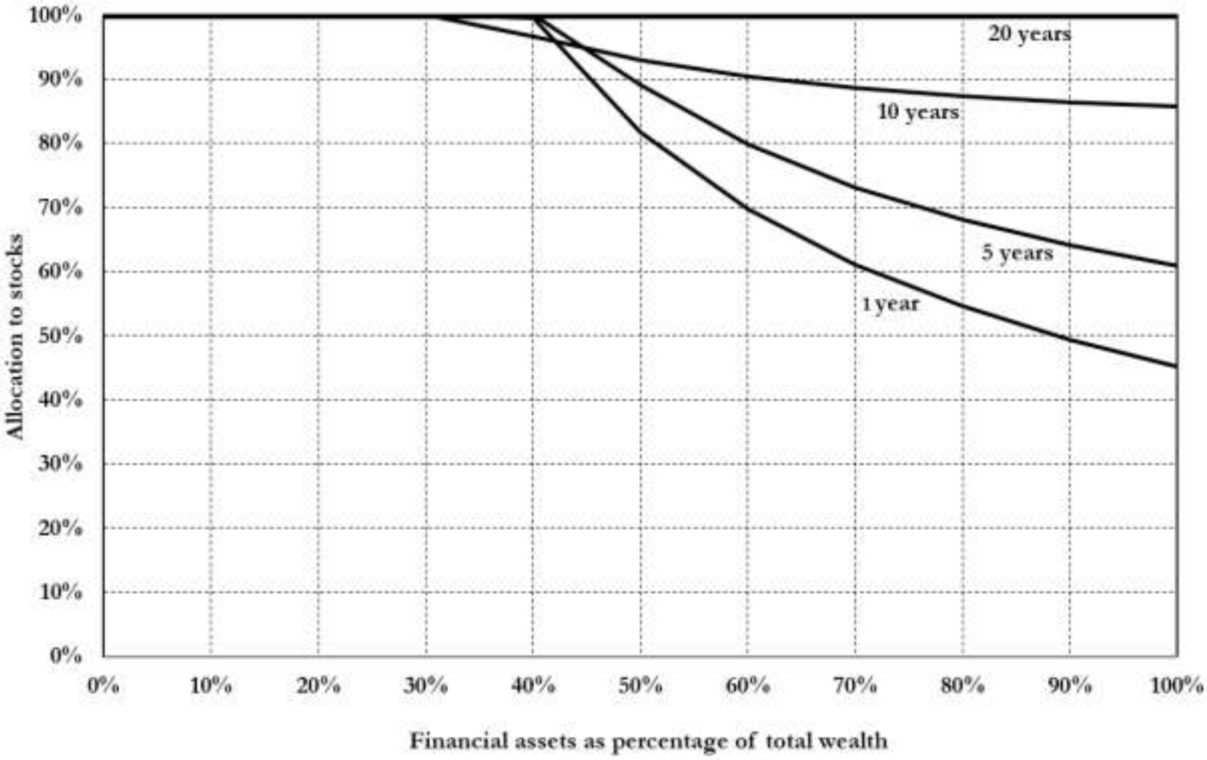

Framework is optimized over different investment horizons using CRRA utility with a risk aversion coefficient of four. The graphs depict stock allocations for 1, 5, 10 and 20-year investment horizons dependent on the percentage of financial wealth as part of total wealth. Real value of non-financial assets is assumed to be fixed during the investment horizon. Real returns are used.

Figure 8 represents an extension of the main expected utility framework presented in figure 5. The figure depicts optimal allocation to stocks in a portfolio given different percentages of financial assets as part of total assets. Each of the recommended allocations for the four investment horizons are for an investor with CRRA utility and a risk aversion coefficient of four. For the 20-year horizon allocation, we can see that the allocation to equities is 100 percent regardless of how much financial assets make up of total wealth.

The graphs for a 1, 5 and 10-year investment horizon all show a downward slope. This means that the recommended equity allocation for a CRRA investor with a risk aversion coefficient of four decreases with the relative weight of financial assets to total assets. These results are not surprising since the existence of fixed non-financial assets represent a risk-free asset. With the assumption of fixed real value, it essentially fixes a certain percentage of the investor's wealth to a risk-free asset which, ceteris paribus, naturally increases the allocation of financial assets to stocks. It is reasonable to assume that an agent who owns a home and has a certain level of human capital will invest his financial portfolio more aggressively compared to someone who has most or all of his wealth tied to his financial portfolio.

The fixed non-financial asset could represent real estate assets, human capital or a combination of both. We make a relatively simple assumption that the non-financial asset is fixed in real value during the agent's investment horizon. This assumption is in part made to make the calculations suitable to our framework, but also because we find a fixed real value to be more realistic for shorter investment horizons. The allocations for the 10 and 20-year horizons should be interpreted with caution. Also, if the value of the non-fixed asset were a combination of real estate and human capital, we would expect the first of these assets to appreciate during the lifetime of the agent, while the other is expected to depreciate. Due to this, one could expect the value development of these two assets to mitigate one another. Lastly, we should point out that this assumption is in line with the assumptions made by Hanna and Chen (1997). They justify their use of a fixed non-financial asset with the fact that the distribution of human wealth rates of return are not available, and that it is very difficult to create reliable long-term estimates of real returns from other asset categories.

One thing that stands out is that both the 1-year line and the 5-year line cross the 10-year line between 40 and 50 percent of financial assets as part of total wealth. This anomaly stand in contrast to time diversification since for this specific area in the figure, an investor with a relatively longer investment horizon is recommended to allocate a smaller part of his or her portfolio to equities than an investor with a relatively shorter investment horizon. However, apart from this isolated anomaly, which we consider to be statistically insignificant, the output speaks in favour of historical time diversification.

The relaxation of Samuelson's third assumption leads to equity recommendations that are more aggressive than the ones in figure 5. We find that the relaxation of the third assumption increases the time diversification effect. The results in figure 8 are similar to the results presented by Hanna and Chen (1997). They also find that the introduction of a fixed non-financial asset increases the optimal equity allocation, however, the equity allocations given by their framework are slightly less aggressive.

5.6 Robustness of equity allocations to changes in the equity risk premium

One risk with using a data set covering such a large part of financial history is that return characteristics may have changed over time, leading to bias in our results. For example, the mean return of stocks may have been considerably higher during the 19th century, and even if the paradigm has changed since then, the higher mean returns will still affect our models.

Optimally, for our models to be valid, we would like the return characteristics to be relatively stable over time. In order to check if this is the case we divide the data into different sub periods.

Table 8 – Mean and standard deviation of real returns for each quartile

	Stocks	Bills	Bonds	Gold	Cash
First quartile, 1802-1855					
Mean	7.32%	5.80%	5.93%	0.76%	0.64%
Std.dev	14.84%	7.33%	7.85%	6.68%	6.68%
Second quartile, 1856-1909					
Mean	9.96%	4.38%	4.47%	0.15%	0.18%
Std.dev	18.36%	5.01%	6.20%	4.41%	4.92%
Third quartile, 1910-1963					
Mean	8.77%	0.12%	1.58%	-0.79%	-2.02%
Std.dev	21.86%	5.92%	8.49%	10.23%	5.62%
Fourth quartile, 1964-2016					
Mean	7.16%	0.92%	3.63%	5.05%	-3.74%
Std.dev	16.91%	2.32%	11.44%	23.11%	2.64%

Table 8 include the mean real returns and standard deviation of real returns for the five asset classes, divided into quartiles.[53] Observing the mean real return of stocks, we can see that these are relatively similar in all four periods. The mean real return of equities in the first and fourth quartile only differ by 0.16 percentage points, which tells us that using data from so far back in time does not bias our framework from the standpoint of mean real return of stocks. However, comparing Bills and gold in the first and fourth quarter we can see that the difference in mean real returns is more striking. Although, since the mean real return of both of these two asset classes have changed over time, we are unable to say whether the most recent characteristics represent a new pattern. Looking at table 7, we cannot find any factor that gives us reason to believe that the optimal equity allocation given by our framework is driven by characteristics in historical data that are no longer prevalent.

However, observing the mean real returns of the assets classes by themselves does not provide us with sufficient information. Optimally, we would like to get a picture of the development of the equity risk premium (ERP).

[53] Summary statistics for the full period can be found in section 4. For a more comprehensive version of table 7, see table A2 in the appendix.

Table 9 – Equity risk premium, for the full period and quartiles

Years	Full period, 1802-2106	First quartile, 1802-1855	Second quartile, 1856-1909	Third quartile, 1910-1963	Fourth quartile, 1964-2016
ERP	4.40%	1.39%	5.48%	7.18%	3.54%

Table 8 shows the ERP for the full period as well as for each of the quartiles. The equity risk premium is the excess return that investing in the stock market provides over the risk-free rate. In our calculations, we have used the return of Bonds as a proxy for the risk-free rate. We could also have chosen Bills as our benchmark since both Bills and Bonds are U.S. Treasury securities, and hence considered risk-free. However, as we can see by looking at table 7 and figure 1, the real return of Bonds and Bills started to diverge, primarily in the 1980s. Ever since this divergence, Bonds have provided a greater real return than Bills. Because of this, we use Bonds as benchmark. We believe it is the more attractive of the two assets from the standpoint of our framework, and hence serve as the better competing asset to equities.

Looking at the ERP for the full period in table 8, we can see that is has been at an average of 4.40 percentage points. This is relatively close, only differing by 0.86 percentage points, to the ERP in the fourth quartile, which is at 3.54 percentage points. The 4.4 percentage points are also quite similar the result obtained by Blanchett et al. (2013), they find an historical ERP of 3.56 percentage points in their 20 country data set. However, looking at the second and third quartile, we can see that the ERP is greater (the ERP for the second and third quartile are on 5.48 and 7.18 percentage points respectively) than the ERP for the fourth quartile. The ERP for the first quartile is in turn considerably smaller than the ERP for the other quartiles, at 1.39 percentage points.

Table 8 tells us that the ERP has not been consistent over time. If the ERP were to change drastically in the future, this could have a considerable effect to a long-term investor allocating assets based our framework. We would therefore like to know how resistant our horizon-based optimal equity allocations are to potential changes in the future ERP.

Table 10 – ERP decrease that would make a CRRA investor with a relative risk coefficient of four indifferent between the strategy depicted in the table and a one-year strategy

Years	5	10	15	20
ERP	-1.24%	-2.09%	-2.55%	-2.60%

Table 9 shows the decrease in the ERP that would make an investor with a CRRA utility function and a risk coefficient of four indifferent between an allocation strategy for 5, 10, 15 or 20-years and an allocation strategy made for one year. To put it differently, if an investor faces a choice between a one-year strategy with a relatively low equity allocation, and a more than one-year strategy with a relatively high equity allocation; how much of a decrease in the ERP would make the strategy with the higher equity allocation provide the same utility as the strategy with the lower equity allocation. For an investor who may be sceptical of the merits of time diversification, this table tells him or her that if you are investing for 10 years, you should allocate assets based on a 10-year time horizon rather than allocating as if you were investing for one year. If the investor were to choose the 10-year strategy, he or she would be able to absorb a 2.09 percentage point decrease in the ERP, and still not obtain less utility than from the allocation of the one-year strategy.

A lower ERP in the future is not implausible. One example of this would be if time diversification were to become accepted by all market participants, and as a result the benefits from it would be fully priced in. In that case, the price of equities would increase relative to other assets, the real returns from equities would decrease relative to risk-free assets, which would lead to a lower ERP. Table 9 tells us that a long-term investor would be able to absorb a considerable degree of "pricing in" of time diversification before the benefits of a time diversification strategy are fully mitigated.

Having found that equity allocations are robust to potential future changes in the ERP, we would like to look more at time diversification. By looking at table 8 we find that the ERP for the fourth quartile is lower than the ERP of the second and third quartile. We would like to make sure that the historical time diversification effect is not driven exclusively by these middle periods. We therefore move on to look at time diversification for only the last of these periods.

5.7 Persistence of time diversification

Figure 9 – Optimal stock allocations, data for the full period versus data for the last 50 years

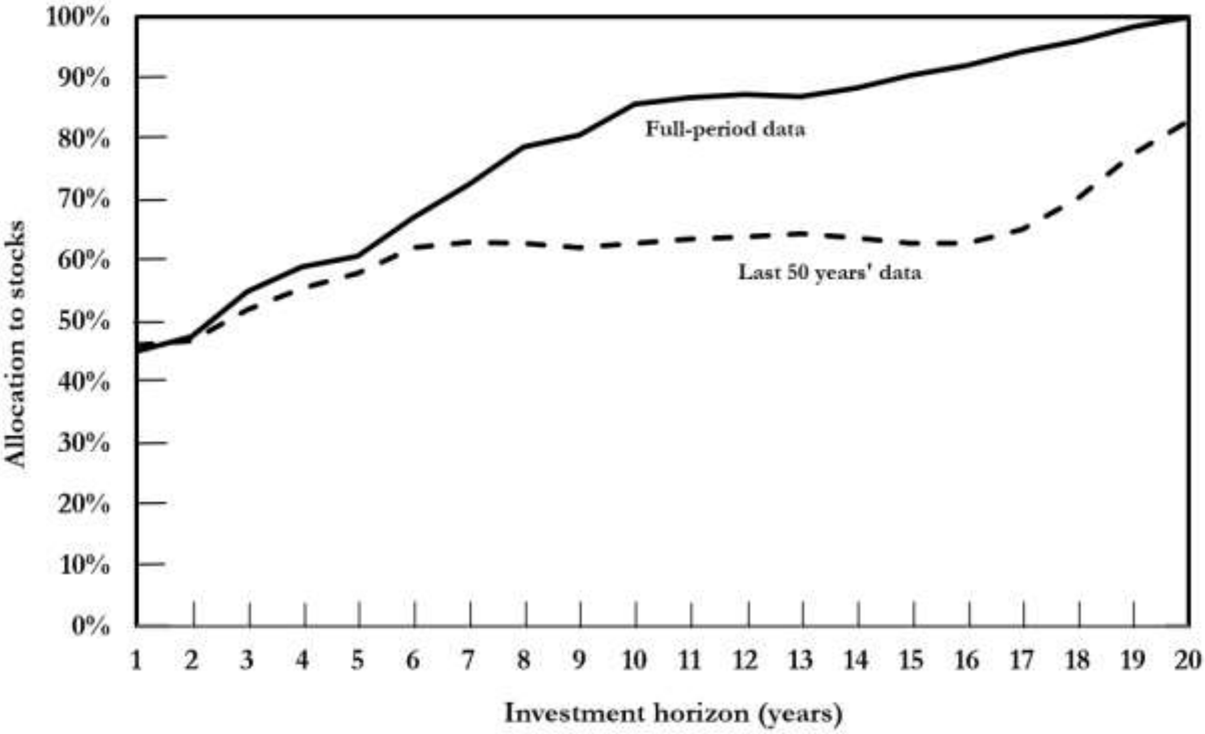

Framework is optimized over investment horizons ranging from 1 to 20 years, with one-year increments. The graphs depict stock allocations for CRRA utility with a risk aversion coefficient of four, when using data for the full-period versus the last 50 years. Real returns are used.

Figure 9 shows the optimal allocation to equities for an investor with CRRA utility and a risk aversion coefficient of four. The solid line depicts the optimal allocation when optimizing using the whole data set, while the dotted line show the optimal equity allocation when optimizing over the last 50 years, i.e. 1966-2016.

By observing figure 9 we can see that the time diversification effect is still prevalent when limiting the optimization to the last 50 years. This gives us indication to believe that it is still rational for a long horizon investor to engage in a time diversification strategy, even if a paradigm shift has occurred 50 years ago which makes the returns before that point in time irrelevant. However, the graphs in figure 9 tell us that the effect of time diversification is less strong for the last 50 years than for the whole period. The two graphs basically show the same equity allocation for an investment horizon of one and two years, while they start to deviate after two years. At a six-year investment horizon we observe that the optimal allocation stops increasing for the dotted line, while the allocation continues to increase for the solid line. In fact, the optimal equity allocations from an investment horizon of six years up until an investment horizon of 16 years are essentially the same. The equity allocation given by the dotted line starts to increase more aggressively after a 16-year investment horizon. In aggregate, the time diversification effect is stronger for the whole period than for the last 50 years, even though the effect is prevalent in both graphs.

We should note that we have chosen not to optimize using more limited data, such as data restricted to the last 20 years, since the conclusions one could draw from such an output would be limited due to the relatively low number of observations. For example, the optimization over a 15-year investment horizon would only occur five times when optimizing using data for 20 years. In addition, we have no justification for claiming that more recent returns are more representative for future returns than returns from further in the past.

Having found that the time diversification effect has not diminished in more recent history, we would like to see if the effect, as well as the previous anomalies in the form of kinks, are still present when using simulated data.

5.8 Simulated data

We are interested in the robustness of time diversification and the existence of the kinks that stand out in figure 3, 4, and 5. Because of this we simulate new returns for the 215 years in our data set in order to see how the optimal allocations changes and if kinks are still present.

The data is simulated using the AR(5) model described in section 5.2. By applying the AR(5) model to the real life historical returns we find a standard deviation of the error term equal to 0.1836. After this we simulate 215 years of new stock returns by using the AR(5) model and a normally distributed error term. The error term is calculated using an expected value of zero and a standard deviation equal to the one we derived from the real life historical returns, i.e. 0.1836.

Figure 10 – Optimal stock allocation for real life and different sets of simulated data, CRRA four

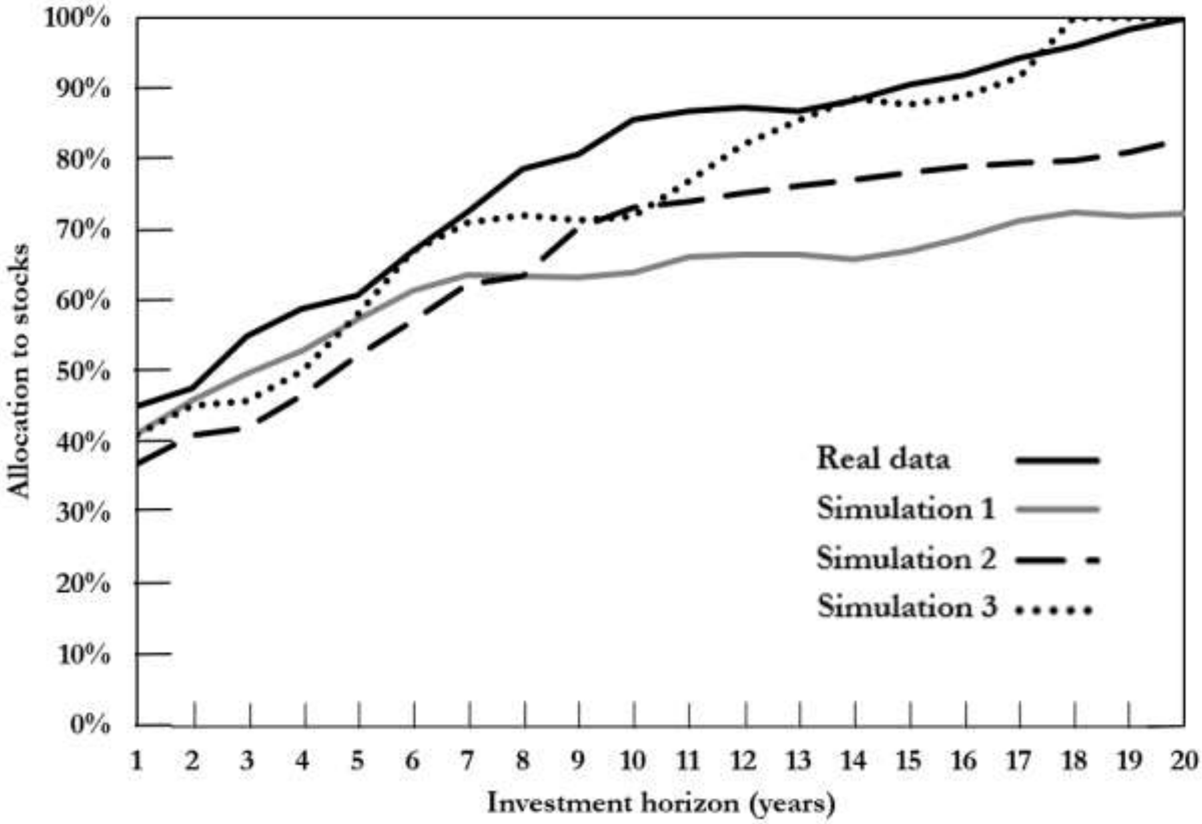

Framework is optimized over investment horizons ranging from 1 to 20 years, with one-year increments. The actual return data is used together with three simulated return series. The graphs depict stock allocations for CRRA utility with a risk aversion coefficient of four. Real returns are used.

Figure 10 depicts the optimal equity allocations for an investor with CRRA utility and a risk aversion coefficient of four, optimized using real life and different sets of simulated data. We can see that in all three simulations, the time diversification effect is still present, but to varying degrees. The real data give the highest optimal equity allocation at all investment horizons apart from 18 and 19 years. However, we believe that this is simply a coincidence due to the limited number of simulations. We should add that we considered a Monte Carlo approach, i.e. running several thousand simulations and showing a graph representing the average of all allocations from the simulations, contrasted against the graph depicting the allocations for the real data. The reason we chose not to do this was that we would have to optimize our framework manually for each time horizon in each simulated dataset. Taking the human aspect as well as the time required by the computer into account, such an approach was considered outside the scope of our project.

Figure 10 shows that time diversification exists also when using simulated data. This is not surprising since we simulate the new data using features from the real data. Unfortunately, figure 10 does not give us any useful information with respect to the kinks in figure 3, 4, and 5. This is partly because figure 10 depicts the allocation for a CRRA investor with a risk aversion of four, which is a profile that have not exhibited any unexpected kinks in the stock allocation for the different investment horizons. Looking at figure 5, we can see that the graphs with significant kinks are the ones depicting risk aversions of eight and above. Because of this we simulate data for a CRRA investor with a risk aversion coefficient of eight.

Figure 11 – Optimal stock allocation for real life and different sets of simulated data, CRRA eight

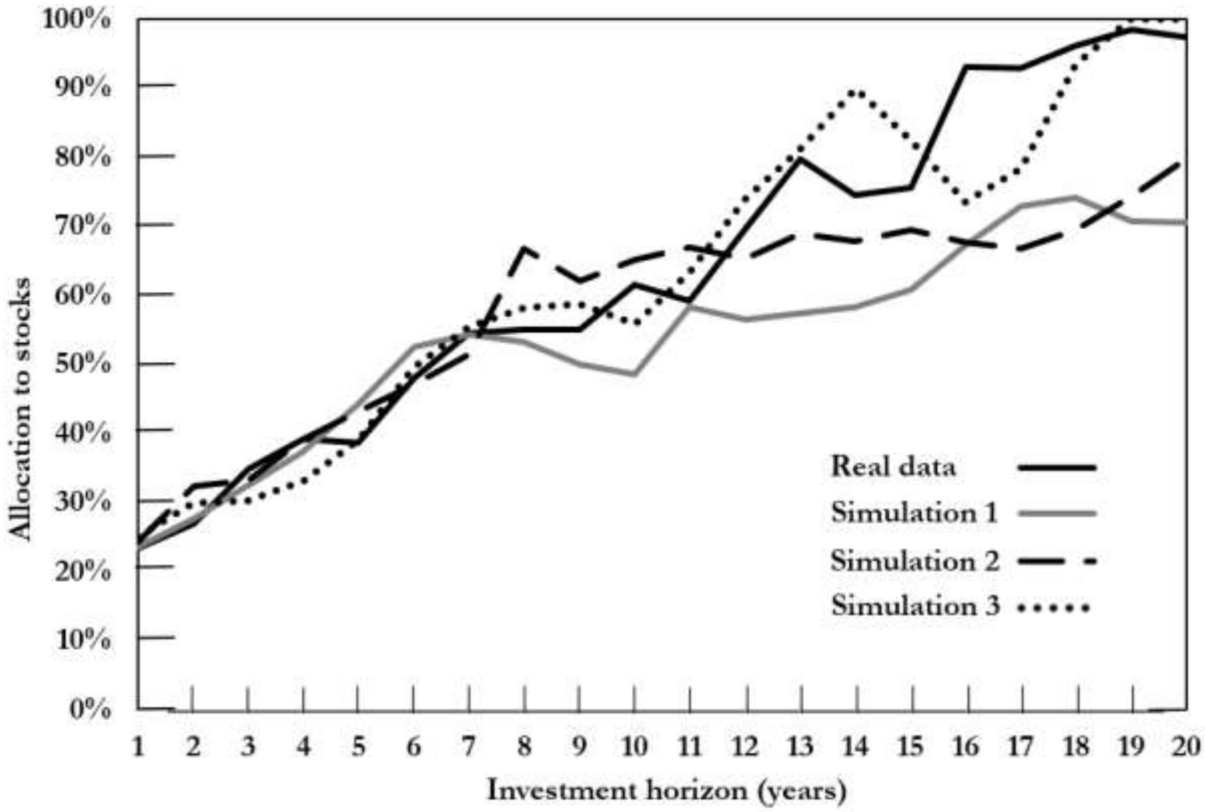

Framework is optimized over investment horizons ranging from 1 to 20 years, with one-year increments. The actual return data is used together with three simulated return series. The graphs depict stock allocations for CRRA utility with a risk aversion coefficient of eight. Real returns are used.

Observing figure 11, we see the different optimal asset allocation for a CRRA investor with a risk aversion coefficient of eight. As in figure 10, three of the four graphs representing the different optimal allocation are all optimized using simulated data. The fourth graph, i.e. the black solid line, represents the optimal allocation using actual historical data.

Looking at the four different graphs we can see that each display kinks at different investment horizons. The dotted line is even, between investment horizons of 13 and 19 years, somewhat inverse to the solid line depicting the actual data. This tells us that if the stock returns follow an AR(5) process, the graphs depicting optimal allocations tend to display kinks, at least for higher levels of risk aversion. The occurrences of kinks seem to be random throughout the investment horizons.

This section has started looking at the kinks in the allocations for an expected utility investor, with the goal of understanding the kinks displayed in figure 5. Next, we look into kinks more thoroughly, also with respect to the anomalies in figure 3 and 4.

5.9 The presence of kinks

This section is devoted to examining and understanding the kinks that are present in some of the previous outputs. Two specific cases are investigated to get a better understanding of the issue; one for the VaR approach and one for the utility approach. It should be noted that given the optimization issues that are dealt with, the explanations for the kinks are most likely due to multiple characteristics of the data. In the following analyses, we simply try to find some of the main drivers. Analyses like the ones we apply here could be done for all single kinks, but we dive into only two selected ones and expect that the remaining kinks have similar causes.

5.9.1 Upward kink in VaR output

First, we will look closer at the upward kink exhibited by the 99 percent certainty -15 percent VaR graph which can be seen in figure 4. Given the 99 percent certainty level we know that just one percent of the outcomes can significantly affect the optimal allocations. We approach the issue by noting that the upward kink is at the three-year horizon; so we look to identify return characteristics in the data that can explain why stocks would be favoured relatively more for the three year horizon as compared to the two and four year horizons. For all three horizons, we have just above 200 observations. This means that the one percent lowest return is (rounded up) the third lowest return. In other words, when we optimize with the VaR, only two out of all the potential outcomes are allowed to be below the VaR limit. Therefore, we look at the third lowest return from the different investment horizons.

Table 11 – Return and years for the third lowest return among the investment horizons, two, three and four years

	Investment horizon		
	2 years	3 years	4 years
Nominal return	-41.46%	-38.10%	-42.94%
Years	1973-1974	2000-2002	1930-1933

Observing table 11, we find supporting evidence of the kink in the data. That is, the third lowest return for a three-year horizon is better than for two and four years. This results in a kink in the optimal allocation making stocks relatively more favourable for a three-year horizon. The specific years where the third lowest returns appear are somewhat equally spread across the full period. In turn, it seems to be arbitrary, and we have no reason to believe that the same pattern will be repeated in the future. An important remark is that when the stock allocation goes down, it specifically means that another asset is included instead. In other words, when examining kinks the stock returns cannot be viewed independently of the alternatives. Looking at the derived optimal portfolios, we see that Bonds are the primary alternative investment. As can be seen from the Bonds characteristics in table 3, Bonds are much less likely than stocks to provide large negative returns, which in turn serve to violate the VaR restriction. We acknowledge that this analysis is not perfect, given that inclusion of the other available assets would be ideal. However, it does provide insight into why and how the kinks come about.

5.9.2 Upward kink in utility output

We go on by looking closer at the downward kink between the investment horizon of 13 and 16 years for the CRRA utility framework with a risk-aversion coefficient of 16.[54] We note that given the behaviour of the utility function, stocks will always be preferred in periods where stocks outperform alternative assets. Therefore, we identify what alternative assets are used when the allocation to stocks decreases. It can be seen from the optimal portfolios that Bills is the main asset class that increases in allocation when Stock allocation falls. Hence, we investigate the periods where bills perform better than stocks. We limit the

[54] See figure 5 for this kink.

analysis to look at the investment horizons ranging from 13 to 16 years to isolate the issue. Ideally, we look for characteristics in the data that can explain a lower (higher) allocation to stocks (bills) for the 14 and 15-year horizon. It is worth reiterating that the coefficient of 16 means that the investor is extremely risk-averse and for that reason is very sensitive to low returns.

Figure 12 – The excess performance of Bills over stocks, limited to periods where Bills outperform stocks

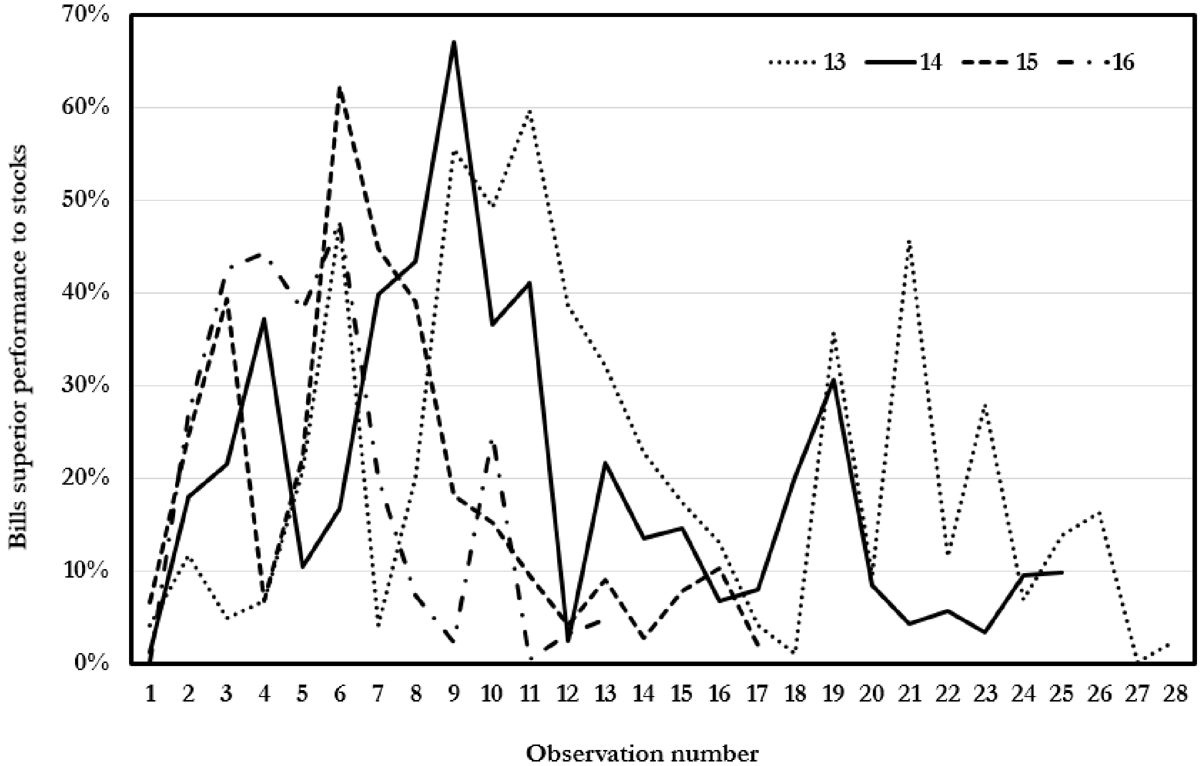

The graphs depict bills excess performance to stocks in percentage points. Real returns are used. The observation numbers refer to instances where bills outperform stocks.

Figure 12 graphs the percentage points excess return that bills provide over stocks, limited to the periods where Bills actually do provide a higher return than stocks. The different graphs shown are for the investment horizons of 13, 14, 15, and 16 years. While the output looks confusing at first, we will focus on the peaks of the different graphs. What can be seen is that the graphs for the 14 and 15-year horizons have the highest peaks. In other words, there are single periods for the 14 and 15-year horizons, where bills excess performance to stocks are higher than in any case for the 13 and 16 horizons. For the interested reader, we note that the 14-year horizon peak is from 1846 to 1859 and the 15-year horizon peak is from 1845-1859. Given the extreme risk-aversion exhibited by the investor, we believe these small details exhibited by the return data have a high impact on the output, i.e. result in kink(s). One again, we find the the specific years to be arbitrary, and will not expect a similar pattern in the future.

We acknowledge several explanations could potentially support the kink but believe that the above data variation plays an important role. To give an example of the complexity, we highlight that the number of times bills outperform stocks changes for the different horizons as can be seen from the different lengths of the graphs depicted in figure 13. Notably it decreases from the 13 to 14-year horizon, which in itself would be an argument for preferring stocks rather than bills.

6. Conclusion and discussion

This paper has optimized portfolios over different investment horizons for agents with different levels of risk preferences, using a dataset covering U.S. return data for number of asset classes for the years 1802–2016. Using a genetic optimization algorithm together with a value at risk (VaR) and an expected utility framework, we found that the optimal allocation to equities tends to increase with the length of the investment period. In addition, the annualized standard deviation of equity returns decreases with time horizon. Our results are in line with Hanna and Chen (1997), Blanchett et al. (2013), Hansson and Persson (2000) and Panyagometh (2011).

We found mean reversion in the real return of equities in our dataset. This meant that our dataset hereby violated Samuelson's second time diversification assumption: the assumption that asset returns are independently and identically distributed (IID). The mean reversion characteristic is sufficiently strong that a moderate relaxation of Samuelson's first and third assumption – the assumptions that the investor is characterized by constant relative risk aversion (CRRA), and that wealth is only a function of returns from financial assets – did not change the fact that the allocation to equities in our optimal portfolios increased with time. Specifically, we were expecting the increasing relative risk aversion (IRRA) to lead to a decreasing or somewhat unchanging equity allocation. However, in relative terms, and in line with Kritzman (1992), extensions to the primary CRRA utility function, in the form of decreasing relative risk aversion (DRRA) showed a stronger time diversification effect than the CRRA, while the opposite was the case when extending the function so that it exhibited IRRA. Similar to Hanna and Chen (1997), the introduction of a fixed non-financial asset leads to more aggressive equity allocations, ceteris paribus.

The effect of time diversification exists when optimizing over the full period, as well as when optimizing over the last 50 years. This implies that our results are not exclusively driven by return features from more than 50 years back. However, the time diversification effect is weaker for the last 50 years in the dataset.

The instability of the optimal allocations, in the form of kinks, can be explained by a limited number of observations in our dataset. Considerable sensitivity to negative returns for the more risk averse investment profiles, together with the fact that some investment periods stand out by showing large negative returns for equities, lead to kinks based on these investment periods. The optimal allocation for agents with high levels of risk aversion can be greatly impacted by only a few such periods with large negative returns.

Our results support the notion that a higher allocation to equities is optimal for agents investing over longer time horizons, thus adding to the empirical evidence supporting time diversification. An investor investing today, with the aim of investing over a longer investment period, should, assuming that the future will look somewhat similar to the past, allocate a relatively larger part of his or her portfolio to equities. Even the most risk averse investors should according to our framework assign 100 percent of their portfolios to equites for long horizons. However, the recommended allocation to risky assets is generally less aggressive than the allocations proposed by Dolvin et al. (2010) and Barberis (2000).

The results generally meet our expectations. Even though Thorley (1995) only showed weak signs of time diversification when optimizing over historical data, and only for investment horizons up to four years, Hanna and Chen (1997), Blanchett et al. (2013) and Siegel (2014) presented more convincing results showing an increased equity allocation with time. However, we were surprised by a number of things, mainly that our main applications, the VaR and CRRA expected utility, gave us very aggressive equity allocations for longer horizons, even for the most risk averse agents. The fact that the mean reversion characteristic in U.S. return data was strong enough to lead to such high equity allocations in our primary outputs, as well as leading to increasing equity allocations in our IRRA extension, was noticeable. The kinks shown in the optimal allocations for more risk averse investors was more prevalent than those found by other scholars such as Blanchett et al. (2013). This could have to do with Blanchett et al. (2013) using a slightly different optimization procedure, but it more likely has to do with them using a data set covering several countries.

The main limitation of the practitioner approach that we have taken, i.e. violating Samuelson's second assumption by optimizing equity allocations based on historical return data, is that we are unable to guarantee that the future will look like the past. No one can be certain that a paradigm shift changing the return process within financial markets will not occur. There is only one sample of financial history and inferences based on a single sample must never be accorded definite interpretations. However, even though this is true, we believe that optimizations based on past returns, or as a second alternative, return processes based on past returns (such as the simulation that is non-IID and based on an autoregressive model, used in this paper), constitute the most plausible approaches. We would argue that the use of simulated returns that are IID and perfectly random, while assuming a constant risk-free rate, often depicted in the form of a binomial tree, does not take into account the realities of investing.[55] Mainly because perfectly random returns include a non-negligible possibility of extremely bad return outcomes, e.g. that the stock market exhibits double digit negative returns for several years in a row. In the case of such extreme outcomes it would not be unreasonable to assume that the risk-free asset is no longer risk-free, thus leading to the original framework becoming invalid. Adopting a framework that assumes a risk-free asset under such conditions could be likened to buying either nuclear insurance or credit default swaps on German Bunds.

Another limitation, relating to Samuelson's first assumption, relates to the loss of practical significance that is inherent in utility functions. Like Kritzman (1994), Levy and Spector (1996), Van Eaton and Conover (1998) as well as Samuelson (1971 and 1994), we acknowledge that our framework could show different results with alternative utility specifications. Rather than exhibiting a continuous utility function, such as the ones applied in this paper, the investor might have a discontinuous utility function. For example, if the agent's wealth was to drop below a certain point this could lead to divorce and having to move to another residence. After having reached that point, the agent would care relatively little about further decreases in wealth. Another example could constitute a reversed scenario, where the agent has become wealthy enough to acquire a mansion and a private jet. Such an individual could be considered relatively indifferent to increases in wealth beyond what he already has. However, as mentioned, our optimization application, due to historical mean reversion, showed increasing equity allocation with

[55] For an example of returns that are perfectly random and IID, depicted in a binomial tree, see table 2 and 3 in Kritzman (1994).

investment horizon for all applied frameworks. And even if another plausible utility function would give us considerably different optimal allocations, it would not be possible to claim that this utility function is the one that best describes the average investor. In addition, the VaR framework is similar to a discontinuous utility function in the sense that the agent can be interpreted to extract infinite negative utility if the VaR level is broken.

Concerning Samuelson's third assumption, we believe it is not realistic to assume that terminal wealth only depends on financial investment performance. We believe our extension including a fixed non-financial asset is more applicable to a real-world scenario than our main utility output. However, a limitation to the application is that it is relatively simplistic in the sense that it only includes a fixed non-financial asset which does not serve as an ideal representation of human capital. A representation which is more detailed and dynamic could come closer to a real-life scenario. For many younger individuals, about to start their career and begin saving, human capital is the most valuable of their assets. In addition, the inclusion of a more suited representation of human capital – through wages or other forms of compensation – means that the agent has the option of adjusting his or her work habits or occupation. For example, if a risky investment was to perform poorly in the beginning or middle of the investment period, the agent could mitigate this effect by working more hours. In that case, longer investment horizons would, ceteris paribus, lead to a higher allocation to equities, since there would be more time available to lessen the negative wealth effects of potential downturns. Having said this, according to us, the problem with such a representation of human capital, and why almost no scholars apart from Pástor and Stambaugh (2012) include it in their framework, is because it makes the model prone to overfitting. A model that is too specific would not be compatible with a practitioner approach. One can also discuss the realism in assuming the agent is able to fully adjust his or her work hours. These factors are often regulated by law and long-term contracts, even if the agent would like to work more to increase his or her income, it is far from certain that this option is available.

Another limitation to our approach is that all optimal allocations concern initial allocations, i.e. the approach does not allow for continuous adjustments over the investment period. This type of non-adjustable allocation is not completely realistic since an investor could, for example, save and invest on a monthly basis. Furthermore, as Detemple (1986) points out, observed state variables in the economy changes the opportunity set over time, which could have an effect on the optimal equity allocation over the investment period. The equity allocation in the optimal portfolios in this paper will also change over the investment horizon due to capital gains and dividends. In addition, allowing for continuous adjustments would most likely lead to overfitted recommendations. The output could for example, depending on the framework, give a recommendation to invest a certain percentage in equities the first year of a five-year horizon, while the percentage for the second year could depend on the returns in the first year. It is not unreasonable to assume the future return process will look like the past in a more general sense, however, it would not be realistic to assume that the pattern and magnitude of business cycles will be exactly the same.

This paper contains a number of insights for mainly market practitioners, and can be used as a tool among others to help decide on investment recommendations to clients. Not counting with a paradigm shift in the future return process of risky assets, clients who are looking to invest for long horizons should be recommended a more aggressive equity allocation. Apart from Panyagometh (2011), there are no papers that we know of that use a VaR when creating optimal portfolios using historical data. In addition, Panyagometh (2011) uses a very limited data set from Thailand which makes his findings hard to extrapolate to more general situations. We believe this output is particularly useful to an actor such as Nordea Wealth Management since the framework corresponds to the classes into which they divide their clients. A client who owns real estate could be relevant to the output showing the changes in optimal allocation when introducing a fixed non-risky asset. Such a client could receive recommendations, all else being equal, to allocate more aggressively to equites than a client who rents his or her residence. Clients could also be interested in seeing how resistant a long-term allocation is to future decreases in the equity risk premium (ERP). Specifically, that it would take a considerable degree of pricing in with respect to the time diversification feature before the client's long-term allocation would equalize the utility of a short-term allocation.

Suggestions to future research would be to try to completely remove the kinks in the graphs for optimal allocations by simulating thousands of returns via a Monte Carlo approach. Such an approach would have to include an automated optimization procedure over every horizon for every simulation. The simulation could be based on a return process derived from historical returns. Future papers could also use alternative data sets covering other nations such as Denmark or the United Kingdom. The main challenge would be to obtain data covering a large enough number of years. Lastly, it would be interesting to see a more detailed incorporation of human capital that is still relatively generalizable and hence useful to a market practitioner. Such an approach could possibly make use of machine learning combined with decision trees in order to decide on allocations based on numerous factors, such as years of education, centre of learning, degree type and years in the labour market.

7. References

Artzner, Philippe, Delbaen, Freddy, Eber, Jean-Marc & Heath, David (1998), "Coherent Measures of Risk", *Mathematical Finance* 9 (3).

Barberis, Nicholas (2000), "Investing for the Long Run When Returns Are Predictable", *The Journal of Finance* 55 (1).

Bennyhoff, Donald G. (2009), "Time diversification and horizon-based asset allocations." *The Journal of Investing* 18 (1).

Bernoulli, Daniel (1738), "Specimen Theoriae Novae de Mensura Sortis", *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5; reprinted in English as "Exposition of a New Theory on the Measurement of Risk", *Econometrica* 22 (1) (1954).

Bianchi, Robert J., Drew, Michael E. & Walk, Adam N. (2016), "The Time Diversification Puzzle: A Survey", *Financial Planning Research Journal* 2 (2).

Blanchett, David, Finke, Michael & Pfau Wade D. (2013), "Optimal Portfolios for the Long Run", *Social Science Research Network*. <ssrn.com/abstract=2320828>.

Blanchett, David, Michael S. Finke & Wade D. Pfau (2013) "Optimal portfolios for the long run." *Social Science Research Network*

Bodie, Zvi (1995), "On the Risk of Stocks in the Long Run", *Financial Analysts Journal* 51 (3).

Brennan M.J. (1997), "The Role of Learning in Dynamic Portfolio Decisions", *European Finance Review* 1 (3).

Butler, Kirt C. & Domian, Dale L. (1991), "Risk, Diversification and the Investment Horizon", *The Journal of Portfolio Management* 17 (3).

Cagan, Philip (1972), "Common Stock Values and Inflation – The Historical Record of Many Countries", *Supplement to NBER Report Thirteen*.

Cagetti, Mark (2003), "Wealth Accumulation Over the Life Cycle and Precautionary Savings", *Journal of Business & Economic Statistics* 21 (3).

Campbell, John Y., Viceira, Luis M. (2002), *Strategic Asset Allocation, Portfolio Choice for Long-Term Investors.* New York: Oxford University Press.

Chinneck, John W (2006), *Practical optimization: a gentle introduction.* <http://www.sce.carleton.ca/faculty/chinneck/po.html>.

Detemple, Jérôme B. (1986), "Asset Pricing in a Production Economy with Incomplete Information", *The Journal of Finance* 41 (2).

Dolvin, Steven D., Templeton, William K. & Rieber, William (2010), "Asset Allocation for Retirement: Simple Heuristics and Target-Date Funds", Journal of Financial Planning 23.

Frontline Systems (2017) (a), "Excel Solver – Algorithms and Methods Used", < http://www.solver.com/excel-solver-algorithms-and-methods-used>.

Frontline Systems (2017) (b), "Excel Solver – GRG Nonlinear Solving Method Stopping Conditions", <http://www.solver.com/excel-solver-grg-nonlinear-solving-method-stopping-conditions>.

Gollier, Christian (2002), “Time Diversification, Liquidity Constraints, and Decreasing Aversion to Risk on Wealth”, *Journal of Monetary Economics* 49 (7).

Hanna, Sherman & Chen, Peng (1997), “Subjective and Objective Risk Tolerance: Implications for Optimal Portfolios”, *Financial Counseling and Planning* 8 (2).

Hansson, Björn & Persson, Mattias (2000), “Time Diversification and Estimation Risk”, *Financial Analysts Journal* 56 (5).

Holland, John H. (1992), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence.* Cambridge: MIT Press.

Janeček, Karel (2004), “What is a Realistic Aversion to Risk for Real-World Individual Investors?” <http://www.sba21.com/personal/RiskAversion.pdf>.

J.P. Morgan (2017), “2017 Long-Term Capital Market Assumptions” <https://am.jpmorgan.com/us/institutional/our-thinking/ltcma-2017>.

Jorion, Philippe (2007), *Value at Risk*. Third edition. New York: McGraw-Hill.

Kritzman, Mark (1992), “What Practitioners Need to Know about Utility”, *Financial Analysts Journal* 48 (3).

Kritzman, Mark (1994), “What Practitioners Need to Know about Time Diversification”, *Financial Analysts Journal* 50 (1).

Kritzman, Mark & Rich, Don (1998), “Beware of Dogma”, *The Journal of Portfolio Management* 24 (4).

Leibowitz, Martin L. & Stanley Kogelman (1991), "Asset Allocation Under Shortfall Constraints." *The Journal of Portfolio Management* 17 (2).

Lee, Hong-Tau, Sheu-Hua Chen & He-Yau Kang (2004), "A Study of Generalized Reduced Gradient Method with Different Search Directions." *Journal of Quantitative Management* 1 (1).

Levy, Haim & Spector, Yishay (1996), “Cross-Asset versus Time Diversification”, *The Journal of Portfolio Management* 22 (3).

Markowitz, Harry (1952), "Portfolio selection." *The Journal of Finance* 7 (1).

Melanie, Mitchell (1999), *An Introduction to Genetic Algorithms.* Fifth printing. Cambridge: MIT Press.

Merrill, Craig & Thorley, Steven (1997), “‘Time Diversification and Option Pricing Theory: Another Perspective’: Response”, *The Journal of Portfolio Management* 23 (4).

Merton, Robert C. (1969), “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case”, *The Review of Economics and Statistics* 51 (3).

Merton, Robert C. & Samuelson, Paul A. (1974), “Fallacy of the Log-Normal Approximation to Portfolio Decision-Making over Many Periods”, *Journal of Financial Economics* 1 (1).

Organisation for Economic Co-operation and Development (OECD) (2015), “Annual Survey of Large Pension Funds and Public Pension Reserve Funds”, < http://www.oecd.org/finance/private-pensions/survey-large-pension-funds.htm>.

Panyagometh, Kamphol (2011), "Asset allocation, time diversification and portfolio optimization for retirement." *Technology and Investment* 2 (2).

Pástor, Ľuboš & Robert F. Stambaugh (2012), "Are Stocks Really Less Volatile in the Long Run?" *The Journal of Finance* 67 (2).

Pérignon, Christophe & Daniel R. Smith (2010), "The Level and Quality of Value-at-Risk Disclosure by Commercial Banks." *Journal of Banking & Finance* 34 (2).

Pritsker, Matthew (2006), "The Hidden Dangers of Historical Simulation." *Journal of Banking & Finance* 30 (2).

Reichenstein, William & Dovalee, Dorsett (1995), *Time Diversification Revisited.* Charlottesville: The Research Foundation of the Institute of Chartered Financial Analysts.

Roy, A. D. (1952), "Safety First and the Holding of Assets", *The Econometric Society* 20 (3).

Samuelson, Paul A. (1963), "Risk and Uncertainty, A Fallacy of Large Numbers", *Scientia* 57 (98).

Samuelson, Paul A. (1969), "Portfolio Selection by Dynamic Stochastic Programming", *The Review of Economics and Statistics* 51 (3).

Samuelson, Paul A. (1971), "Maximizing the Geometric Mean in Long Sequences of Investing or Gambling", *The Journal of Investing* 19 (4).

Samuelson, Paul A. (1989), "The Judgment of Economic Science on Rational Portfolio Management: Indexing, Timing, and Long-Horizon Effects", *The Journal of Portfolio Management* 16 (1).

Samuelson, Paul A. (1994), "The Long-Term Case for Equities", *The Journal of Portfolio Management* 21 (1).

Stangeland, David A. & Turtle, Harry J. (1999), "Time Diversification: Fact or Fallacy", *Journal of Financial Education* 25.

Siegel, Jeremy J. (2014), *Stocks for the Long Run: The Definitive Guide to Financial Market Returns & Long-Term Investment Strategies (EBOOK)*. Fifth edition. New York: McGraw Hill Education.

Siegel, Jeremy J. (2016), "Historical Data", <http://www.jeremysiegel.com/index.cfm/Fuseaction/Resources.ListResources&type=data.cfm>.

Sowell, Thomas (2014), *Basic Economics.* Fifth edition. New York: Basic Books.

Thorley, Steven R. (1995), "The Time-Diversification Controversy", *Financial Analysts Journal* 51 (3).

Van Eaton, R. Douglas & Conover, James A. (1998), "Misconceptions about Optimal Equity Allocation and Investment Horizon", *Financial Analysts Journal* 54 (2).

Stock, James H. & Watson, Mark (2012), *Introduction to Econometrics: International Edition*. Third edition E. London: Pearson Higher Education.

Willis Towers Watson (2017), "Global Pension Assets Study 2017", <https://www.willistowerswatson.com/en/insights/2017/01/global-pensions-asset-study-2017>.

8. Appendix

Derivation of equation 15 based on equation 13

$$U(W) = \frac{W^{1-\alpha}}{1-\alpha} \tag{A1}$$

$$U'(W) = (1-\alpha) \cdot \frac{W^{-\alpha}}{1-\alpha} = W^{-\alpha} \tag{A2}$$

$$U''(W) = -\alpha \cdot W^{-\alpha-1} \tag{A3}$$

$$RRA(W) = -W \cdot \frac{U''(W)}{U'(W)} = -W \cdot \frac{-\alpha \cdot W^{-\alpha-1}}{W^{-\alpha}} = \alpha \tag{A4}$$

Derivation of equation 15 based on equation 14

$$U(W) = \ln(W) \tag{A5}$$

$$U'(W) = \frac{1}{W} \tag{A6}$$

$$U''(W) = -W^{-2} \tag{A7}$$

$$RRA(W) = -W \cdot \frac{U''(W)}{U'(W)} = -W \cdot \frac{-W^{-2}}{W^{-1}} = 1 \tag{A8}$$

Derivation of equation 18 based on equation 16

$$U(W) = \frac{(W-\theta)^{1-\alpha}}{1-\alpha} \tag{A9}$$

$$U'(W) = (1-\alpha) \cdot \frac{(W-\theta)^{-\alpha}}{1-\alpha} = (W-\theta)^{-\alpha} \tag{A10}$$

$$U''(W) = -\alpha \cdot (W-\theta)^{-\alpha-1} \tag{A11}$$

$$RRA(W) = -W \cdot \frac{U''(W)}{U'(W)} = -W \cdot \frac{-\alpha(W-\theta)^{-\alpha-1}}{(W-\theta)^{-\alpha}} = \frac{\alpha \cdot W}{W-\theta} = \frac{\alpha}{1-\theta/W} \tag{A12}$$

$$For\ \theta = 0 \rightarrow RRA(W) = \alpha \tag{A13}$$

Derivation of equation 18 based on equation 17

$$U(W) = \ln(W-\theta) \tag{A14}$$

$$U'(W) = \frac{1}{W-\theta} \tag{A15}$$

$$U''(W) = -(W-\theta)^{-2} \tag{A16}$$

$$RRA(W) = -W \cdot \frac{U''(W)}{U'(W)} = -W \cdot \frac{-(W-\theta)^{-2}}{(W-\theta)^{-1}} = \frac{W}{W-\theta} = \frac{1}{1-\theta/W} \tag{A17}$$

$$For\ \theta = 0 \rightarrow RRA(W) = 1 \tag{A18}$$

Derivation of absolute risk aversion (ARA) applied on the utility function described by equation 16 and 17

$$U(W) = \frac{(W-\theta)^{1-\alpha}}{1-\alpha} \tag{A19}$$

$$U'(W) = (1-\alpha) \cdot \frac{(W-\theta)^{-\alpha}}{1-\alpha} = (W-\theta)^{-\alpha} \tag{A20}$$

$$U''(W) = -\alpha \cdot (W-\theta)^{-\alpha-1} \tag{A21}$$

$$ARA(W) = \frac{-U''(W)}{U'(W)} = \frac{-\alpha(W-\theta)^{-\alpha-1}}{(W-\theta)^{-\alpha}} = \frac{\alpha}{W-\theta} \tag{A22}$$

Figure A1 – Optimal stock allocations for different VaR-levels, 90 percent certainty

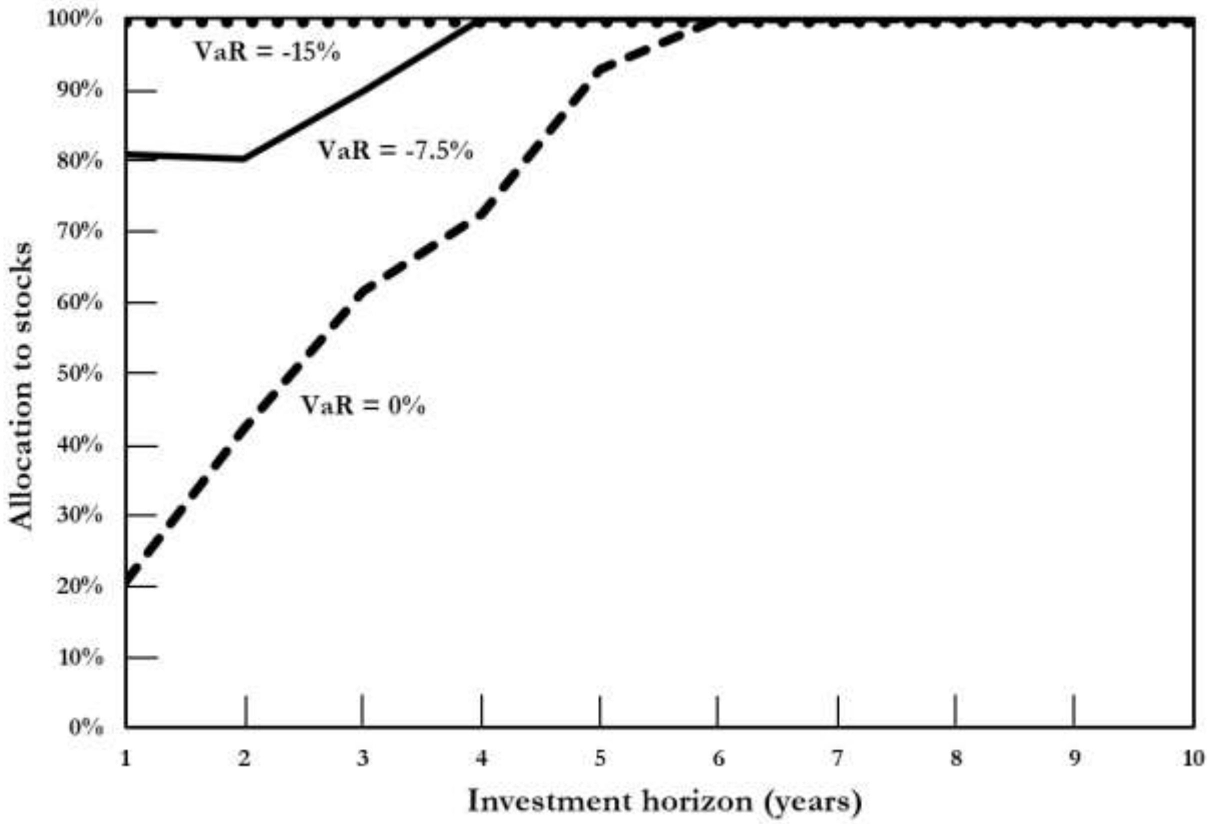

Framework is optimized over investment horizons ranging from 1 to 10 years, with one-year increments. Five different VaR-levels are used, -30 percent, -20 percent, -15 percent, -7.5 percent and 0 percent. Each graph depicts the optimal allocation for a specific VaR-level. Graphs for -20 and -30 percent are not shown because they overlap perfectly with the one for -15 percent. Nominal returns are used.

Table A1 – SAS results from AR(2) model

		Parameter Estimates		
Variable	Estimate	Standard Error	t Value	Pr > |t|
Intercept	0.0961	0.0146	6.60	<0.0001
Lag1	0.0103	0.0679	0.15	0.8797
Lag2	-0.1734	0.0676	-2.57	0.0110
		Ordinary Least Squares Estimates		
MSE	0.0319		DFE	210
MAE	0.1369		Root MSE	0.1787
MAPE	199.8085		AIC	-126.0920
ESS	0.2110		TSS	6.879
SSE	6.7078		R-Square	0.0305

Table A2 – Summary statistics of real returns for each quartile

	Stocks	Bills	Bonds	Gold	Cash
		First quartile, 1802-1855			
Mean	7.32%	5.80%	5.93%	0.76%	0.64%
Std.dev	14.84%	7.33%	7.85%	6.68%	6.68%
Median	6.38%	5.45%	6.42%	0.21%	0.21%
Maximum	62.38%	23.68%	29.76%	17.98%	17.98%
Minimum	-29.87%	-9.59%	-9.92%	-13.67%	-13.67%
Skewness	0.7652	0.4014	0.4020	0.3104	0.3182
Kurtosis	2.8071	0.1993	0.5945	0.3345	0.3685
		Second quartile, 1856-1909			
	Stocks	Bills	Bonds	Gold	Cash
Mean	9.96%	4.38%	4.47%	0.15%	0.18%
Std.dev	18.36%	5.01%	6.20%	4.41%	4.92%
Median	6.92%	4.17%	4.52%	0.22%	0.52%
Maximum	66.62%	12.28%	15.46%	10.23%	7.69%
Minimum	-28.85%	-15.63%	-21.86%	-20.95%	-21.28%
Skewness	0.4228	-2.1165	-1.4958	-1.9324	-2.7407
Kurtosis	1.2044	7.2214	5.1201	9.5646	9.8692
		Third quartile, 1910-1963			
	Stocks	Bills	Bonds	Gold	Cash
Mean	8.77%	0.12%	1.58%	-0.79%	-2.02%
Std.dev	21.86%	5.92%	8.49%	10.23%	5.62%
Median	5.77%	1.18%	1.02%	-1.29%	-1.62%
Maximum	57.15%	17.38%	30.26%	60.50%	11.97%
Minimum	-38.57%	-15.07%	-16.91%	-17.00%	-17.00%
Skewness	-0.0167	-0.0657	0.4270	4.0492	-0.3076
Kurtosis	-0.4015	1.4990	1.6013	24.6000	1.9759
		Fourth quartile, 1964-2016			
	Stocks	Bills	Bonds	Gold	Cash
Mean	7.16%	0.92%	3.63%	5.05%	-3.74%
Std.dev	16.91%	2.32%	11.44%	23.11%	2.64%
Median	8.71%	1.85%	3.22%	-0.62%	-1.48%
Maximum	57.15%	17.38%	35.13%	99.94%	11.97%
Minimum	-37.29%	-3.74%	-14.54%	-38.13%	-11.74%
Skewness	-0.7589	0.3188	0.5687	1.4804	-1.4222
Kurtosis	0.1350	-0.2990	0.0089	4.6513	1.9881

Table A3 – Optimal portfolios for zero percent VaR with 95 percent confidence

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	0%	0%	0%	0%	100%
2	29%	19%	52%	0%	0%
3	46%	22%	30%	2%	0%
4	57%	23%	9%	11%	0%
5	79%	9%	10%	2%	0%
6	66%	15%	17%	2%	0%
7	87%	12%	1%	0%	0%
8	96%	3%	1%	0%	0%
9	100%	0%	0%	0%	0%
10	100%	0%	0%	0%	0%

Table A4 – Optimal portfolios for -7.5 percent VaR with 95 percent confidence

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	53%	11%	28%	7%	0%
2	59%	8%	33%	0%	0%
3	74%	0%	26%	0%	0%
4	83%	0%	10%	6%	0%
5	92%	0%	8%	0%	0%
6	81%	10%	10%	0%	0%
7	100%	0%	0%	0%	0%
8	100%	0%	0%	0%	0%
9	100%	0%	0%	0%	0%
10	100%	0%	0%	0%	0%

Table A5 – Optimal portfolios for -15 percent VaR with 95 percent confidence

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	83%	5%	11%	1%	0%
2	87%	7%	6%	0%	0%
3	99%	0%	0%	1%	0%
4	100%	0%	0%	0%	0%
5	95%	0%	5%	0%	0%
6	100%	0%	0%	0%	0%
7	100%	0%	0%	0%	0%
8	100%	0%	0%	0%	0%
9	100%	0%	0%	0%	0%
10	100%	0%	0%	0%	0%

Table A6 – Optimal portfolios for -20 percent VaR with 95 percent confidence

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	100%	0%	0%	0%	0%
2	100%	0%	0%	0%	0%
3	100%	0%	0%	0%	0%
4	100%	0%	0%	0%	0%
5	100%	0%	0%	0%	0%
6	100%	0%	0%	0%	0%
7	100%	0%	0%	0%	0%
8	100%	0%	0%	0%	0%
9	100%	0%	0%	0%	0%
10	100%	0%	0%	0%	0%

Table A7 – Optimal portfolios for -30 percent VaR with 95 percent confidence

Investment horizon	Stocks	Bonds	Bills	Gold	Cash
1	100%	0%	0%	0%	0%
2	100%	0%	0%	0%	0%
3	100%	0%	0%	0%	0%
4	100%	0%	0%	0%	0%
5	100%	0%	0%	0%	0%
6	100%	0%	0%	0%	0%
7	100%	0%	0%	0%	0%
8	100%	0%	0%	0%	0%
9	100%	0%	0%	0%	0%
10	100%	0%	0%	0%	0%

Table A8 – Optimal portfolios for CRRA with a coefficient of one

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	100%	0%	0%	0%	0%
2	100%	0%	0%	0%	0%
3	100%	0%	0%	0%	0%
4	100%	0%	0%	0%	0%
5	100%	0%	0%	0%	0%
6	100%	0%	0%	0%	0%
7	100%	0%	0%	0%	0%
8	100%	0%	0%	0%	0%
9	100%	0%	0%	0%	0%
10	100%	0%	0%	0%	0%
11	100%	0%	0%	0%	0%
12	100%	0%	0%	0%	0%
13	100%	0%	0%	0%	0%
14	100%	0%	0%	0%	0%
15	100%	0%	0%	0%	0%
16	100%	0%	0%	0%	0%
17	100%	0%	0%	0%	0%
18	100%	0%	0%	0%	0%
19	100%	0%	0%	0%	0%
20	100%	0%	0%	0%	0%

Table A9 – Optimal portfolios for CRRA with a coefficient of two

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	81%	0%	19%	0%	0%
2	81%	0%	19%	0%	0%
3	86%	0%	14%	0%	0%
4	89%	0%	11%	0%	0%
5	94%	0%	6%	1%	0%
6	97%	0%	0%	3%	0%
7	97%	0%	0%	3%	0%
8	97%	0%	0%	3%	0%
9	96%	0%	0%	4%	0%
10	96%	0%	0%	4%	0%
11	97%	0%	0%	3%	0%
12	97%	0%	0%	3%	0%
13	99%	0%	0%	1%	0%
14	100%	0%	0%	0%	0%
15	100%	0%	0%	0%	0%
16	100%	0%	0%	0%	0%
17	100%	0%	0%	0%	0%
18	100%	0%	0%	0%	0%
19	100%	0%	0%	0%	0%
20	100%	0%	0%	0%	0%

Table A10 – Optimal portfolios for CRRA with a coefficient of four

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	45%	3%	52%	0%	0%
2	48%	0%	46%	6%	0%
3	55%	0%	39%	6%	0%
4	59%	0%	35%	6%	0%
5	61%	0%	32%	7%	0%
6	67%	0%	22%	11%	0%
7	73%	0%	16%	11%	0%
8	79%	0%	9%	12%	0%
9	81%	0%	6%	13%	0%
10	86%	0%	0%	14%	0%
11	87%	0%	0%	13%	0%
12	87%	0%	0%	13%	0%
13	87%	0%	0%	13%	0%
14	88%	0%	0%	12%	0%
15	91%	0%	0%	9%	0%
16	92%	0%	0%	8%	0%
17	94%	0%	0%	6%	0%
18	96%	0%	0%	4%	0%
19	98%	0%	0%	2%	0%
20	100%	0%	0%	0%	0%

Table A11 – Optimal portfolios for CRRA with a coefficient of eight

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	23%	51%	21%	4%	0%
2	27%	51%	14%	8%	0%
3	35%	48%	10%	6%	0%
4	39%	48%	7%	5%	0%
5	39%	50%	6%	6%	0%
6	48%	35%	8%	9%	0%
7	55%	32%	5%	9%	0%
8	55%	37%	0%	8%	0%
9	55%	36%	0%	9%	0%
10	62%	29%	0%	10%	0%
11	59%	33%	0%	8%	0%
12	70%	22%	0%	9%	0%
13	80%	7%	1%	12%	0%
14	75%	17%	0%	9%	0%
15	76%	18%	0%	6%	0%
16	93%	0%	0%	7%	0%
17	93%	0%	0%	7%	0%
18	96%	0%	0%	4%	0%
19	98%	0%	0%	2%	0%
20	97%	0%	0%	3%	0%

Table A12 – Optimal portfolios for CRRA with a coefficient of 16

Investment horizon	Stocks	Bills	Bonds	Gold	Cash
1	12%	76%	2%	10%	0%
2	12%	75%	0%	13%	0%
3	21%	79%	0%	0%	0%
4	26%	74%	0%	0%	0%
5	20%	80%	0%	0%	0%
6	41%	58%	0%	2%	0%
7	42%	57%	0%	2%	0%
8	35%	64%	0%	0%	0%
9	38%	61%	0%	1%	0%
10	40%	59%	0%	1%	0%
11	40%	59%	0%	1%	0%
12	58%	39%	0%	2%	0%
13	69%	25%	0%	6%	0%
14	43%	55%	0%	2%	0%
15	39%	60%	0%	1%	0%
16	87%	10%	0%	2%	0%
17	95%	0%	0%	5%	0%
18	100%	0%	0%	0%	0%
19	100%	0%	0%	0%	0%
20	100%	0%	0%	0%	0%