Network for Studies on Pensions, Aging and Retirement

NETSPAR ACADEMIC SERIES

Rebalancing for long-term investors

Joost Driessen, Ivo Kuiper

DP 05/2017-013

Rebalancing for long-term investors

Joost Driessen*, Ivo Kuiper†

October 12, 2017

Abstract

In this study we show that the rebalance frequency of a multi-asset portfolio has only limited impact on the utility of a long-term passive investor. Although continuous rebalancing is optimal, the loss of a suboptimal strategy corresponds to up to only 30 basis points of the initial wealth of the investor, assuming market returns are unpredictable and transaction costs can be ignored. Our results suggest that reducing transaction costs clearly outweighs the benefit of frequent rebalancing. When we study a setting where asset returns are predictable, we find that a long-term investor that ignores this predictability underestimates the benefit of less frequent rebalancing. In this setting, limiting the frequency to at least once every quarter results in significant higher utility, even without transaction costs.

1 Introduction

An increasing amount of investors consider themselves to be passive investors. In this study we focus specifically on long-term investors with multi-asset portfolios who believe that the investment opportunity set is constant, or at least they have no (access to) valuable insights in changes of this set. For these investors, for example a pension fund or endowment, rebalancing trades are a large part of portfolio turnover and resulting trading costs can be significant. They often consider a fixed asset allocation to be the best way to achieve their investment goals. Over time, the portfolio weights drift away from the established model portfolio. After all, the weight of a category in the portfolio changes due to the relative price developments. Rebalancing is often seen as essential maintenance of a portfolio, in order to keep the expected risk and return properties of the portfolio within boundaries. By doing this the investor maximizes the expected utility. As the total market is by definition buy and hold, not every investor can rebalance as explained by Sharpe (2010). Setting the rebalancing strategy is therefore one of the few active decisions these investors face besides setting their target portfolio. Although our framework can also be used to evaluate rebalancing within an asset class, we focus on the impact of rebalancing on the strategic portfolio level of asset classes as this level is the

*Tilburg University, Department of Finance; CentER Tilburg University.
†Kempen Capital Management. Corresponding author: ivo.kuiper@kempen.nl.

most important driver of portfolio returns.

In this study we quantify the utility costs of periodic rebalancing of portfolios that consist of multiple risky assets, in our case equity and government bonds. Compared to Constantinides (1986) and Balduzzi and Lynch (1999), who study the optimal rebalancing strategy including transaction costs, we suggest a complementary approach. This allows us to evaluate multi-asset portfolios with non-IID returns in contrast to portfolios containing a single risky asset. Increasing the number of risky assets makes it hard to find the optimal solution, especially when the investor is not myopic and dynamic programming is needed. In this study we limit ourselves to myopic investors and focus on the impact of the rebalancing frequency in multiple environments.

Our main finding is that, although continuous rebalancing is optimal, the utility cost of less frequent rebalancing when market returns are unpredictable is limited to only basis points in terms of the initial wealth of the investor over a twenty year period. For a simplified case we show analytically that not rebalancing leads to both drift and variation in portfolio weights. This results in an increasing utility cost of the multi-asset portfolio when the rebalancing period grows. Economically the impact is small. This cost turns into a gain if we take transaction costs into account. Our results suggest that reducing transaction costs by less frequent rebalancing clearly outweigh optimizing the portfolio more frequently, unless transaction costs are extremely small. Using simulations, we show that these results hold for different levels of risk aversion and expected returns. Our calculations show that, when taking transaction costs into account, most of the marginal gain is captured when limiting the rebalancing frequency to once every year.

When we study a setting where asset returns are predictable, we find that a long-term investor that ignores this predictability underestimates the benefit of less frequent rebalancing. If expected market returns are time-varying, there is an additional utility effect as a result of trends or mean reversion. Dependent on the time series dynamics, this effect can be positive or negative. Using a VAR model to model the predictability, and ignoring transaction costs, we estimate that the utility gain of rebalancing once every year compared to once every month is about 1.5 percent in terms of the investors' initial wealth over a twenty-year period. Hence by less frequent rebalancing the investor profits from trending behavior of market returns. About half this benefit is already captured by limiting the rebalancing frequency to once every quarter instead of monthly.

In the next section, section 2, we discuss the methodology used. Section 3 is about the impact of rebalancing when market returns are unpredictable. In 4 we look at the impact of return predictability. Conclusions and suggestions are the subject of section 5.

2 Methodology

A rebalancing policy makes sure that the portfolio composition remains within limits, but it also protect investors against several behavioural biases. In an environment with a lot of volatility, rebalancing has the most impact[1]. Immediately after a stock market crash, there are few investors who want to step into the market. A portfolio manager who automatically rebalances does. The same is true the other way around. Rebalancing helps to prevent a pro-cyclical policy[2]. Sometimes it is suggested that rebalancing results in a more diversified portfolio than a buy and hold strategy. That argument, however, assumes that the model portfolio offers optimal diversification. That does not have to be right as several market theories assume that the market portfolio is the most diversified portfolio (no idiosyncratic risk). When rebalancing implies larger deviations from the market portfolio, it would actually result in a decrease in diversification[3].

Our starting point is that every investor wants to maximize its utility. The realized utility is a function of the portfolio chosen and the realized market returns. Every time period an rational investor determines the portfolio weights such that, based on return and risk assumptions and the investors utility function, the expected utility is maximized.

In order to model the impact of rebalancing on our long-term passive investor we assume that our investor has power utility. As described by Campbell and Viceira (2002) power utility implies that the relative risk aversion γ is constant (CRRA), and the absolute risk aversion is declining in wealth. CRRA implies that the initial wealth of an investor does not affect the optimal portfolio. The power utility function is given by

$$\begin{aligned} U(W_{t+K}) &= \frac{W_{t+K}^{1-\gamma} - 1}{1-\gamma}, \\ \lim_{\gamma \to 1} U(W_{t+K}) &= log(W_{t+K}), \end{aligned} \tag{1}$$

where U denotes utility, W stands for wealth and γ is the constant relative risk aversion. In Figure

[1] Ang (2014) shows that rebalancing is short volatility.

[2] A portfolio manager who believes in mean reversion, does this even stronger. Obviously there is a possibility that buying after a crash is not profitable (this time is different). A rebalancing strategy can therefore be considered to be short regime shifts.

[3] In addition, there is often a reference to a better return / risk ratio as result of the rebalancing strategy. Several papers [eg Bouchey et al. (2012) and Erb and Harvey (2006)] suggest that rebalancing yields higher returns, a so-called rebalance premium by looking at the geometric returns. The premium is attributed to rebalancing or diversification, and illustrated with quotes like 'just as it is possible to harness energy from waves in the ocean, it is possible to harvest returns from volatility in the market' by Bouchey et al. (2012) and 'turning water into wine' Erb and Harvey (2006). Chambers et al. (2012) explains that the geometric return is an inferior measure of the expected return, because a higher geometric returns is not a good indicator of long-term expected wealth. Although realized geometric average returns correctly sort realized total returns, this is not the case for geometric expected returns and expected total returns. Maximizing an expected geometric efficiency is, therefore, not an optimal portfolio strategy. Less volatility does not increase the expected wealth, a higher expected arithmetic return does.

5.1 we show utility as function of the expected relative change in wealth for various values of risk aversion. It clearly shows the impact of risk aversion, utility decreases much more rapidly in the case of a loss and benefits less from potential gains.

In this study we use a portfolio with a risk-free asset and two risky assets, i.e. an equity and government bond index. As mentioned, we focus on rebalancing on a strategic level, and do not consider the rebalancing of the underlying stocks in the equity index. Limiting ourselves to three assets enable us to gain useful insights, while not making the analysis overly complex. Our passive investor assumes that the market returns are IID. Earlier studies, Samuelson (1969), Merton (1969) and Merton (1971), already showed that a long-term investor with power utility chooses the same portfolio as a short-term investor if the returns are IID. IID returns are a very reasonable assumption for our passive investor. CRRA implies that wealth, and therefore past returns, do not impact portfolio choice. IID implies that there won't be new information the next period, so therefore the long-term investor holds a constant portfolio equal to the portfolio that a short term investor prefers.

The investment horizon is an important aspect of the utility function and investment strategy. We want our investor's horizon to approximate the horizon of a typical long-term investor. We assume that a horizon of 20 years fits this description. Also Balduzzi and Lynch (1999) look at the terminal wealth at this horizon in their study. With monthly data this implies 239 opportunities to rebalance.

In order to measure the wealth impact of rebalancing, we evaluate two cases of market dynamics. First, we assume log normal IID returns. In this case the assumption of our investor about market dynamics matches reality. Second, we assume that returns are described by a VAR model resulting in a mismatch between the investment beliefs of our investor and the true market dynamics.

Our analysis enables us to quantify the impact of reducing the rebalancing frequency on the utility of our investor in both cases. Some cases we are able to solve analytically, for others we run simulations.

When market returns are IID, the impact of rebalancing is driven by the utility loss as result of suboptimal portfolio weights. In the case of the VAR returns, our results are also driven by the extent to which the allowed drift in the portfolio overlaps with the market dynamics estimated by the VAR model.

Introducing predictability in asset returns impacts the optimal portfolio, as an investor wants to hedge against changes in the opportunity set. As summarized by Cochrane (2001) and Ang (2014), in the case of log utility the income and substitution effects exactly cancel out, removing the intertemporal hedging demands and thereby enabling myopic portfolio choice. The true optimal weights are still time varying, however log utility ensures that the average optimal weights are equal to the optimal weights in the case of unpredictable returns. As a consequence, the difference in utility impact by introducing predictability is solely because of ignoring the time series behavior in markets, and not because of a suboptimal average portfolio composition. Note that a risk aversion coefficient of 1 is relatively low. For example, Balduzzi and Lynch (1999) use a range between 2 and 10. In the IID case, we show that our results also hold for higher levels of risk aversion.

As optimal portfolio we choose a portfolio that is fully invested, without leverage. The reason for this is that if we have a leveraged portfolio this would result in defaults when we do not frequently rebalance, violating our assumption of log normality of returns[4].

2.1 Transaction costs

Transaction costs can have a significant impact on investment returns. In earlier studies different type of costs are used. The most common are fixed cost per trade, costs proportional to amount (fixed percentage) and costs that are quadratic to the traded amount. The last model is the most common way to take market impact into account.

Of course, combinations are also possible. Operating costs such as gathering and processing information are usually fixed, while the paid transaction costs are proportionate (or quadratic to take market impact into account for an illiquid asset). Transaction costs can vary by asset and time (less liquid, and thus higher costs, declining market etc).

Constantinides (1986) calculates the optimal no-trading range for a portfolio of one risky asset and a risk-free asset as a function of proportional transaction costs. Realistic trading costs result in a sizable no trading range. Constantinides (1986) concludes that an additional 0.2% annual return provides enough compensation when the transaction costs are 1.0%. This liquidity premium is there to compensate for the lower demand for the risky asset. Balduzzi and Lynch (1999) use proportional costs of 0.5% in the base case. Bikker et al. (2007) show that average market impact costs equal 20 basis points for equity bought and 26 basis points for equity sold by pension funds. In our study we use simulations to show the impact of transaction costs on our results. For these simulations we assume constant proportional one way trading costs of 0.1% for both the equity and government bond

[4]Based on historical return data, log utility would lead to a portfolio with a hefty leverage of about 300%.

index.

2.2 Rebalancing strategies

There are several common rebalancing strategies. In this study we look at constant-mix strategies (Perold and Sharpe (1995)). In the simplest case the portfolio weights are adjusted towards the target weights every period, regardless of the deviation from this target. This is the optimal choice when transaction costs are ignored. A second method is a method in which a minimum deviation is defined before one rebalances, a no trading range. The portfolio is only rebalanced when the weight is outside these boundaries. Ang (2014) refers to these strategies as 'calender' and 'contingent' rebalancing. Drivers of the trading frequency are thus both exogenous (trading costs) and endogenous (market dynamics). There are various trades possible if the weights of an asset is outside the no trading range. One might adjust the weight at once to the target weight (eg when when only assuming fixed costs). Another approach is to bring the weight back to the limit of the bandwidth (in the case of costs proportional to amount). Finally, there is an approach in which the rebalance is done in as many steps as possible (optimal when assuming market impact, with quadratic costs to amount). In the situation where transaction costs are a function of the price development, asymmetrical bandwidths can be optimal. In the situation that costs are high during and after depreciation, the hurdle to buy is higher than to sell. In addition to the size of the trade, the timing of execution may vary. A distinction can be made between contingent (continuous monitoring) and calendar rebalancing (periodic monitoring). Liu (2004) concludes that for realistic trading costs the optimal trading frequency is about once a year. An important limitation to the results of both Liu (2004) and Constantinides (1986) is the assumption that there is no predictability of market returns, and the opportunity is set constant. In this study we focus on the impact of the rebalancing frequency. Our base case is a periodic rebalance to the optimal weight. This is the optimal strategy if we ignore trading costs.

3 Passive investor with unpredictable returns

In this section we show the impact of rebalancing for a passive investor, assuming that market returns are IID and log normal distributed. Our passive investor has no intention of forecasting the markets. He believes that time variation in future returns cannot be forecasted and in this case we assume that he is right. As a result the investment beliefs of the investor correspond with reality. First we will derive the utility curve for this investor as function of the portfolio weights. Then, we build intuition analytically on what the impact of less rebalancing is, in the case of log utility. As mentioned in the introduction, log utility implies a low risk aversion coefficient ($\gamma = 1$). We use simulations to quantify the impact for the more complex cases, including higher levels of risk aversion ($\gamma>1$). We first discuss the portfolio choice and utility function for a single time period, then we continue with multi period expressions in order to show the impact of rebalancing.

By definition, the utility function specifies what the investors cares about. The goal of every investor is to maximize his utility U, often as function of wealth W. For a single period Campbell and Viceira (2002) and Mulvey and Simsek (2002) summarize this as

$$\max \mathrm{E}_t U(W_{t+1}), \tag{2}$$

subject to the arithmetic return of the portfolio R_p:

$$W_t = (1 + R_{p,t})W_{t-1}. \tag{3}$$

As we assume our investor has log utility[5], we can write utility U for a single period as

$$\begin{aligned} \mathrm{E}_t U\left(\frac{W_{t+1}}{W_t}\right) &= \mathrm{E}_t \log\frac{W_{t+1}}{W_t} \\ &= \mathrm{E}_t \log(1 + R_{p,t+1}) \\ &= \mathrm{E}_t r_{p,t+1}, \end{aligned} \tag{4}$$

where the utility equals the expected log return of the portfolio, $r_{p,t}$. Following the proposed approximation of the log portfolio return by Campbell and Viceira (2002) for the case with multiple risky assets we know that for short time intervals the log portfolio return r_p can be expressed as

$$r_{p,t+1} - r_{f,t+1} = \boldsymbol{\alpha}_t'(\boldsymbol{r}_{t+1} - r_{f,t+1}\boldsymbol{i}) + \frac{1}{2}\boldsymbol{\alpha}_t'\boldsymbol{\sigma}_t^2 - \frac{1}{2}\boldsymbol{\alpha}_t'\boldsymbol{\Sigma}_t\boldsymbol{\alpha}, \tag{5}$$

where $\boldsymbol{\alpha}$ denotes the portfolio weights, $\boldsymbol{r}$ represent a vector containing the log returns of the individual assets, $\boldsymbol{\Sigma}$ is the covariance matrix of which $\boldsymbol{\sigma}^2$ is the diagonal, r_f denotes the risk-free return which is in our case the 1-month US Treasury Bill rate and $\boldsymbol{i}$ is a vector containing ones.

[5]As mentioned, we want our investor to behave myopically, in order to make sure that the optimal portfolio weights are the same for every rebalancing period. In the case of IID returns, log utility is not a requirement for myopic portfolio choice. However, in order to be able to compare the results with a situation without IID returns, we assume log utility also in this case. In section 3.5 we run the same analysis for other levels of risk aversion.

The optimal weights of the portfolio are related to the expected risk premiums of the risky assets. Common is to optimize the portfolio weights based on expected returns. As mentioned, we prefer no leverage in our portfolio as we want a zero probability for default. We therefore choose to calculate the risk premiums that corresponds to the portfolio weights of our preference. In this study we use 60% equity and 40% bonds, $\boldsymbol{\alpha}_t = [0.6, 0.4]$. For each set of $\boldsymbol{\alpha}_t$ the expected returns are given by

$$\mathrm{E}_t \mathbf{r}_{t+1} = r_{f,t+1}\boldsymbol{i} + \gamma \boldsymbol{\Sigma}_t \boldsymbol{\alpha}_t - \boldsymbol{\sigma}_t^2/2. \tag{6}$$

Given a covariance matrix $\boldsymbol{\Sigma}$, one can calculate the expected (annualized) returns. For the equity index we use stock data from the CRSP US stock database including NYSE, NYSE MKT and NASDAQ with dividends. For our government bond returns we use 10 year US Treasury returns. Based on monthly data from 1973 to 2015 we calculate the covariance matrix $\boldsymbol{\Sigma}$,

$$\boldsymbol{\Sigma} = \begin{pmatrix} 0.0021 & 0.0001 \\ 0.0001 & 0.0005 \end{pmatrix}. \tag{7}$$

These numbers correspond to an annual standard deviation of 15.9% for equity and 7.8% for bonds. The average correlation between monthly stock and bond returns is 9.8%.

If we use these expected returns to fit our conditions, we find an expected annualised arithmetic return $\mathrm{E}R$ of 5.9% for the equity index and 4.6% for the bond index, which, given an historical annual risk-free rate of 4.3%, corresponds to an annual risk premium of respectively 1.6% and 0.3%. In section 3.5 we show that higher risk premiums lead to similar results.

Figure 5.2a and 5.2b show the resulting utility of a single period for different weights of the risky assets, in the case of log utility (equation (5)). We clearly see that the utility decreases exponentially with the distance from the optimal weights. However, we also see that the impact of suboptimal weigths on the utility is very limited in absolute terms[6].

3.1 Impact of rebalancing

Until now we have discussed utility for a single time period. As our investor is a long-term investor, it is the utility of investing over a long horizon that matters (20 years, K=240). Assuming log normal

[6]an increase in the equity weight from 60% to 80% requires a equity return of +167% assuming that bonds keep their value, and this corresponds to a 1% decrease in utility over one (very short) timeperiod. In section 3.5 we show the impact of the level of risk aversion γ on the utility curve.

returns and log utility, we can write equation (4) for horizon K as

$$\begin{aligned}
\mathrm{E}_t U\left(\frac{W_{t+K}}{W_t}\right) &= \mathrm{E}_t \log\left(\frac{W_{t+K}}{W_t}\right) \\
&= \mathrm{E}_t \log\left(\frac{W_{t+1}}{W_t}\cdot\frac{W_{t+2}}{W_{t+1}}\cdot\ldots\cdot\frac{W_{t+K}}{W_{t+K-1}}\right) \\
&= \sum_{i=1}^{K} \mathrm{E}_t \log\left(1+R_{p,t+i}\right) \\
&= \sum_{i=1}^{K} \mathrm{E}_t r_{p,t+i} \\
&\approx \sum_{i=1}^{K} r_f + \mathrm{E}_t \boldsymbol{\alpha}'_{t+i-1}\left(\mathrm{E}_t \boldsymbol{r}_{t+i} - r_f \boldsymbol{i}\right) + \frac{1}{2}\mathrm{E}_t \boldsymbol{\alpha}'_{t+i-1}\boldsymbol{\sigma}^2_{t+i-1} - \\
&\quad \frac{1}{2}\mathrm{E}_t\left(\boldsymbol{\alpha}'_{t+i-1}\boldsymbol{\Sigma}_{t+i-1}\boldsymbol{\alpha}_{t+i-1}\right).
\end{aligned} \tag{8}$$

When the investor rebalances every month ($n = 1$), we know that the weights $\boldsymbol{\alpha}_t$ are the same for every t. Therefore, in an IID world with log normal returns, the portfolio has a constant expected log return $\mathrm{E}r_p$. In the case of log utility, the utility is linear with horizon K and can be expressed as

$$\begin{aligned}
\mathrm{E}U_{n1}\left(\frac{W_{t+K}}{W_t}\right) &= K \cdot \mathrm{E}r_p \\
&\approx K \cdot \left(r_f + \boldsymbol{\alpha}'(\mathrm{E}\boldsymbol{r} - r_f\boldsymbol{i}) + \frac{1}{2}\boldsymbol{\alpha}'\boldsymbol{\sigma}^2 - \frac{1}{2}\boldsymbol{\alpha}'\boldsymbol{\Sigma}\boldsymbol{\alpha}\right).
\end{aligned} \tag{9}$$

The expression becomes more complex when the investor does not rebalance every period, as now the weights of the assets in the risky portfolio are no longer constant but stochastic with drift. Suppose our investor rebalances every n months, and K is a multiple of n, the expression (still assuming log normal IID returns and log utility) is given in equation (10) below. For interpretation purposes, we split the sum of the time intervals in a double sum, using the rebalancing period n. The last step makes use of the fact that after a rebalance moment, the expectations of the portfolio weights are the same as after the previous rebalance moment. We can therefore calculate the overall utility as a sum of K/n 'blocks':

$$
\begin{aligned}
\mathrm{E}_t U_n\left(\frac{W_{t+K}}{W_t}\right) &\approx \sum_{i=1}^{K} r_f + \mathrm{E}_t \boldsymbol{\alpha}'_{t+i-1}(\mathrm{E}\boldsymbol{r} - r_f \boldsymbol{i}) + \frac{1}{2}\mathrm{E}_t \boldsymbol{\alpha}'_{t+i-1}\boldsymbol{\sigma}^2 - \\
&\quad \frac{1}{2}\mathrm{E}_t\left(\boldsymbol{\alpha}'_{t+i-1}\boldsymbol{\Sigma}\boldsymbol{\alpha}_{t+i-1}\right) \\
&= \sum_{j=1}^{K/n}\sum_{i=1}^{n} r_f + \mathrm{E}_t \boldsymbol{\alpha}'_{t+(j-1)n+i-1}(\mathrm{E}\boldsymbol{r} - r_f \boldsymbol{i}) + \frac{1}{2}\mathrm{E}_t \boldsymbol{\alpha}'_{t+(j-1)n+i-1}\boldsymbol{\sigma}^2 - \\
&\quad \frac{1}{2}\mathrm{E}_t\left(\boldsymbol{\alpha}'_{t+(j-1)n+i-1}\boldsymbol{\Sigma}\boldsymbol{\alpha}_{t+(j-1)n+i-1}\right) \\
&= \frac{K}{n}\sum_{i=1}^{n} r_f + \mathrm{E}_t \boldsymbol{\alpha}'_{t+i-1}(\mathrm{E}\boldsymbol{r} - r_f \boldsymbol{i}) + \frac{1}{2}\mathrm{E}_t \boldsymbol{\alpha}'_{t+i-1}\boldsymbol{\sigma}^2 - \\
&\quad \frac{1}{2}\mathrm{E}_t\left(\boldsymbol{\alpha}'_{t+i-1}\boldsymbol{\Sigma}\boldsymbol{\alpha}_{t+i-1}\right).
\end{aligned}
\tag{10}
$$

In order to gain insight into the impact of rebalancing, we derive the expression for the simplified case that both K and n are 2 and there is only one risky asset:

$$
\begin{aligned}
\mathrm{E}_t U\left(\frac{W_{t+2}}{W_t}\right) &\approx \sum_{i=1}^{2} r_f + \mathrm{E}_t \alpha_{t+i-1}(\mathrm{E}r - r_f) + \frac{1}{2}\mathrm{E}_t \alpha_{t+i-1}\sigma^2 - \frac{1}{2}\mathrm{E}_t\left(\alpha^2_{t+i-1}\sigma^2\right) \\
&= \underbrace{r_f + \alpha_t(\mathrm{E}r - r_f) + \frac{1}{2}\sigma^2\alpha_t - \frac{1}{2}\sigma^2\alpha_t^2}_{i=1} + \\
&\quad \underbrace{r_f + \mathrm{E}_t\alpha_{t+1}(\mathrm{E}r - r_f) + \frac{1}{2}\sigma^2\mathrm{E}_t\alpha_{t+1} - \frac{1}{2}\sigma^2\left(\left(\mathrm{E}_t\alpha_{t+1}\right)^2 + \mathrm{Var}_t\alpha^2_{t+1}\right)}_{i=2}.
\end{aligned}
\tag{11}
$$

In equation (11) we can distinguish various effects. For the first time interval,($i = 1$) the weights are known, and utility rises when the expected return increases and decreases with portfolio volatility. For the second time interval however ($i = 2$) there is both drift and stochastic uncertainty around the weights at time t. As this illustration shows, these both effects results in an additional drag on utility.

3.2 Results

Because we ignore transaction costs, we know that the optimal rebalancing frequency is $n = 1$ (continuous rebalancing). With increasing n, the portfolio weights become more suboptimal as a result of both drift and the stochastic process, and therefore the expected utility decreases. As visualised in Figure 5.2a and 5.2b, the utility decay is parabolic with the difference between the actual weight and the optimal weight of the risky assets. This leads to a similar non linear effect in the overall utility as function of the rebalancing frequency. In Figure 5.3 we show the results of our simulations for rebalancing every month and rebalancing every two years ($n = 24$). Figure 5.3b shows the distribution of the portfolio weights when $n = 24$. Combining this distribution with the utility function in Figure 5.2b the expected impact of rebalancing is small in this case.

Figure 5.4a shows the expected wealth and confidence interval for both continuous rebalancing ($n = 1$) and rebalancing every 24 months. In both cases the horizon K is 240 months. We find only a slightly higher expected return and confidence intervals in the case of less rebalancing. Figure 5.4b shows the (indexed) utility of our investor as function of the rebalancing frequency n. Panel A in Table 5.2 shows both the absolute and indexed utility. In this case, we find that the impact is very limited.

For every rebalance period the same initial wealth results in a different utility. We can express the loss in utility by calculating the difference in initial wealth needed to compensate for this[7]. In order to do this, we calculate what initial wealth would generate the same utility as in the optimal case, given the expected log wealth. The expression for log utility ($\gamma = 1$) is given in equation (12). U_1 denotes the optimal utility with initial wealth W_0 and final wealth distribution W_1. Furthermore, we have suboptimal utility U_n, wealth distribution W_n and adjusted initial wealth W_0^*.

When we define $U_{n1} = \text{Elog}\left(\frac{W_{n1}}{W_0}\right)$ and $U_n = \text{Elog}\left(\frac{W_n}{W_0}\right)$, it follows that

$$\begin{aligned} U_{n1} &= \text{Elog}\left(\frac{W_n}{W_0^*}\right) \\ &= \text{Elog}\left(W_n\right) - \log W_0 - \log W_0^* + \log W_0 \\ &= U_n + \log\left(\frac{W_0}{W_0^*}\right) \\ e^{(U_{n1} - U_n)} &= \frac{W_0}{W_0^*} \\ W_0^* &= W_0 \cdot e^{(U_n - U_{n1})}. \end{aligned} \tag{12}$$

A similar expression for higher levels of risk aversion is complex. In order to calculate the results for more risk averse investors, as we do to evaluate the robustness of our results, we solve the wealth equivalent numerically. In Figure 5.5 and Panel A of table 5.2 we show the utility costs in terms of initial wealth equivalent. The loss in utility is limited to only basispoints (0.00-0.04%) of the initial wealth W_0 on a horizon K of 20 years. Note that this analysis ignores trading costs. The utility costs increase with the investment horizon K. In section 3.4 we shows that reducing transaction costs outweigh the benefits of the utility gained by frequently rebalancing.

[7] This is very similar to the certainty equivalent approach where the investor evaluates the amount of money to be received with certainty that will make him/her indifferent between the risk-free and the risky cash flows.

3.3 Suboptimal initial portfolio

Until this point, we have assumed that our investor has the optimal portfolio at start. Our results show that the impact of rebalancing is very small in that case. Ignoring transaction costs, there is no meaningful (economic) difference between the expected utility of an investor that rebalances monthly versus one that rebalances only every two years. As explained in section 2, one characteristic of periodic rebalancing is that it does not take into account the actual deviations from the optimal weights. The implicit underlying assumption is that the frequency is high enough to not care about this. However, we are also interested in the impact of outliers. What is the impact when the investor is faced with extreme market conditions, when the next rebalancing moment is still far away? In this section we evaluate the impact when the initial portfolio composition deviates significantly from the optimal portfolio. The sensitivity of the expected utility to the initial portfolio composition shows the costs of waiting until the next rebalancing moment for investors faced with extreme returns short after a rebalancing moment. This analysis provides insight to the benefit of implementing allocation bandwidths as (additional) rebalancing triggers.

We find that the utility costs of waiting until the next rebalance period are higher in this case, although the overall impact is still small in terms of initial wealth. Panel B of Table 5.2 and Figure 5.6 shows the utility as function of the rebalance period, assuming a start portfolio of 70% equity and 30% bonds (while the optimal portfolio is unchanged, 60% and 40%). Compared to the results with the results in Panel A we find only a slightly larger impact of the rebalancing frequency on the utility. The suboptimal initial portfolio leads to somewhat higher uncertainty in the wealth after 240 months. When the initial portfolio has more than the optimal amount of bonds, we find a similar impact on utility. In this case the dominant force behind the decrease by a lower expected return (in contrast to higher uncertainty in the final wealth).

These results give only limited support to the idea that it is smart to set rebalance triggers on the portfolio weight of risky assets instead of periodic rebalancing. Even when the next rebalancing is still 24 months away, the impact is small.

3.4 Impact of transaction costs

As mentioned in the introduction, for passive long-term investors rebalancing trades are often the major part of portfolio turnover, and resulting trading costs can be significant. By definition, trading costs are a hurdle for transactions and reduce the optimal amount of trading. The gain in utility by rebalancing the weights must offset set the loss in utility by making trading costs. We show the results assuming proportional, single way transaction costs of 0.1%. In Figure 5.7a we show that the utility gain is not enough and there is a large loss in expected utility. The difference in utility between continuous rebalancing and once every 24 months corresponds to almost 0.7% of initial wealth, following equation (12). These results are in line with the sizeable no trade region in Constantinides (1986) for a single risky asset.

The total amount of trading costs is in this study assumed to be proportional with the traded volume. Table 5.1 shows the estimated (annualised) turnover for the portfolio for the different rebalance periods, together with the standard deviation of this average in our simulations. Over the investment horizon rebalancing every month results in an annualised traded volume of about 40 percent of the portfolio. This corresponds to about eight times the initial portfolio size over the total 20 years. Trading costs of 0.1% mean that in the case of continuous rebalancing, total trading costs over the investment horizon of 20 years are less than a percent of the initial wealth. Reducing the rebalancing frequency reduces this volume significantly. Other rebalancing strategies, such as the ones including a no trade range, result in lower expected turnover.

3.5 Impact of risk aversion

In the calculations so far we have assumed that the risk aversion γ is equal to one, resulting in log utility. Based on this level of risk aversion, we calculated the risk premiums that matched our optimal portfolio of choice. As higher risk aversion levels are common in practise, we look at the robustness of our results. We find that for more risk averse investors, the impact of rebalancing is also very limited.

A more risk averse investor prefers a different portfolio based on the same risk premiums. While the optimal risky portfolio is the same for every investor for every set of risk premia, the optimal allocation to this risky portfolio decreases proportionally with $1/\gamma$, the inverse of the risk aversion. If an investor with log utility is fully invested, an investor with $\gamma = 2$, invests half of total assets in the risk-free asset. In order to have the same portfolio, a higher risk premium is required. In this section we show the results of these two alternatives. First, we keep the risk premiums constant and show the resulting impact of rebalancing assuming different optimal portfolios. Second, we keep the optimal portfolio constant and use different values for the risk premiums. This enables us to evaluate how sensitive our results are to our assumption about the risk premiums. We find that higher risk premiums result in more utility loss by less frequently rebalancing, roughly a factor 10. However, the impact is still very limited in economic sense.

Figure 5.8 shows the utility curve for different values of risk aversion. Figure 5.8a shows the utility curves, while keeping the optimal portfolio constant to 60% equities and 40% bonds (and different levels of the expected risk premiums). Figure 5.8b shows the curve while keeping the risk premium constant (and as a result different optimal weights for each investor).

Panel A of Table 5.3 and figure 5.9 show that for investors with a higher risk aversion, the utility cost of rebalancing less is reduced when the risk premium is constant, explained by a much smaller allocation to portfolio of risky assets. This effect is larger than the increased aversion to stochastic behaviour of the composition of the risky portfolio. Remember that the risk premiums in our previous calculations are relatively low. These results show that higher risk premiums, combined with a higher level of risk aversion, result in similar levels of utility loss. In terms of initial wealth equivalent however, the impact is a factor 5 to 10 higher.

Panel B of Table 5.3 shows both the optimal portfolios and the impact of rebalancing on expected utility for various values of risk aversion. If an investor invests less, the utility decreases because of lower expected wealth. In this case, rebalancing is less important as volatility of the portfolio is also lower. On the other hand, a risk averse investor is more sensitive for the stochastic behaviour of the portfolio, negative surprises have more impact on utility than positive surprises.

4 Passive investor with predictable returns

In the previous chapter we assumed that our passive investor was right about the market dynamics. In this section we show the results when market returns are not IID but predictable.

An IID distribution implies return series without autocorrelation. When prices are mean reverting the expected return will increase because of rebalancing. When prices are trending the expected return is lowered by the rebalance. The expected autocorrelation of the returns of a category therefore plays an important role in the choice of a rebalance policy. In this section we evaluate the impact of this behaviour on the utility cost of rebalancing. As alternative to IID returns, we assume that the market returns are described by a first order VAR model.

Introducing predictability in asset returns impacts the optimal portfolio. In this case, the analytical solution for the utility over K periods (equation (10)) is no longer valid, but our passive investor believes it is. The true optimal weights are time varying, however the average optimal weight is still

60% equity and 40% bonds as our investor has log utility (no intertemporal hedging demand, also in the case of VAR returns). As a consequence, the difference in utility is solely because of ignoring the time series behavior in markets, and not because of a suboptimal average portfolio composition.

4.1 VAR model

VAR models are dynamic systems of equations in which the current level of each variable in the system depends on past movement in that variable and in all the other variables in the system. VAR models are frequently used in literature to describe market dynamics.

The VAR approach seeks to identify revisions in expectations by using the time-series structure. It postulates that the unobserved components of returns can be written as linear combinations of innovations to observable variables.

In order to calibrate our analysis, we use monthly data series starting in 1973. Once every month is also the highest rebalancing frequency that we consider. Similar to Campbell et al. (2003) we use value weighted equity returns on the NYSE, NASDAQ and AMEX from the Center for Research in Security Prices (CRSP).

For the VAR model, we also use the 10 year bond return from the UST and Inflation Series (CTI) in CRSP and as risk-free rate the 1 month Treasury bill rate. In order to calculate the yieldspread we downloaded the 5 year zero coupon yield from the CRSP Fama Bliss data.

The VAR approach starts by defining a vector z_{t+1} which has k elements. To a large extent we follow the approach used by Campbell et al. (2003). Balduzzi and Lynch (1999) only uses the dividend yield as predictive variable in addition to the asset returns. In our case the first three are the real short rate r_t, the equity return e_t and the 10 year treasury return b_t. Furthermore we use the annual equity return ey_t, the nominal short rate n_t, the dividend price ratio D/P_t, the yieldspread S_t and the CP factor CP_t, introduced by Cochrane and Piazzesi (2005):

$$\boldsymbol{r}_t = [r_t, e_t, b_t, ey_t, n_t, D/P_t, S_t, CP_t]. \tag{13}$$

We use both a monthly and annual equity return in the model in order to better capture both short and longer term dynamics. We now have a vector $\boldsymbol{r}_t$ of p directly observable variables that are to be forecasted as well as indicators that are helpful doing that. In this study we use a VAR model with a single lag:

$$\boldsymbol{r}_{t+1} = \boldsymbol{c} + \boldsymbol{A} \cdot \boldsymbol{r}_t + \boldsymbol{\varepsilon}_{t+1}. \tag{14}$$

We use a single lag because of the increasing risk of overfitting when using a higher order VAR. The variable $\boldsymbol{A}$ denotes the auto regression matrix and $\boldsymbol{\varepsilon}_{t+1}$ is a vector with error term containing the unexpected components with covariance matrix $\boldsymbol{\Sigma}_V$:

$$\boldsymbol{\varepsilon}_{t+1} \sim N(0, \boldsymbol{\Sigma}_V). \tag{15}$$

All of the variables appear to be stationary in our sample period[8]. For simulation purposes, $\boldsymbol{\Sigma}_V$ is different than $\boldsymbol{\Sigma}$ used in the previous section as a result of the time varying expected returns:

$$\begin{aligned} \boldsymbol{\Sigma}_{Vt} &= \mathrm{Cov}_t(\boldsymbol{\varepsilon}_{t+1}) \\ &= \mathrm{Cov}_t(\boldsymbol{r}_{t+1} - \mathrm{E}_t \boldsymbol{r}_{t+1}) \\ &= \mathrm{Cov}_t(\boldsymbol{r}_{t+1} - \boldsymbol{A} \cdot \boldsymbol{r}_t). \end{aligned} \tag{16}$$

The passive investor, ignoring these dynamics, will use $\boldsymbol{\Sigma}$, the same covariance matrix as under the IID assumption[9]. We use the unconditional means as starting values for the simulated returns.

Table 5.4 shows the estimated autoregression matrix $\boldsymbol{A}$. The explanatory power of excess equity returns of our model is 2.7%. This is lower than the 8.6% found in Campbell et al. (2003). This difference is due to calculations on monthly frequency instead of quarterly and, to a lesser extent, the use of the 1m Tbill instead of the 3 month Tbill as risk-free rate and a longer sample.

4.2 Impact of rebalancing

Figure 5.11 and panel A of Table 5.6 show the impact of the rebalance period on the portfolio returns and utility. We find that in the case of predictable returns rebalancing results in larger utility costs than in the case of unpredictable returns. Based on these results, it is beneficial to have a longer rebalance period, even when ignoring transactions costs. About half of the maximum utility gain

[8] In order to be able to compare results with the IID case, we adust e_t and b_t such that ER is equal to the values in section 3.

[9] Our passive investor assumes that the expected returns are constant. This implies $\mathrm{Cov}(\mathrm{E}_t \boldsymbol{r}_{t+1}) = 0$.

$$\begin{aligned} \mathrm{Cov}(\boldsymbol{r}_{t+1}) &= \boldsymbol{\Sigma} = \mathrm{E}\mathrm{Cov}_t(\boldsymbol{r}_{t+1}) + \mathrm{Cov}(\mathrm{E}_t \boldsymbol{r}_{t+1}) \\ &= \boldsymbol{\Sigma}_V + \mathrm{Cov}(\boldsymbol{A} \cdot \boldsymbol{r}_t) \\ &= \boldsymbol{\Sigma}_V + \boldsymbol{A} \boldsymbol{\Sigma}_V \boldsymbol{A}' \end{aligned} \tag{17}$$

compared to continuous rebalancing is achieved when rebalancing every three months. This large utility increase for quarterly and yearly rebalancing in stead of monthly rebalancing can be explained by trending returns. Increasing the rebalance period further results in only small marginal utility gains. In terms of initial wealth equivalent, we find that when an investor rebalances every 24 months, this investor should start with about 1.5% less wealth in order to have the same expected utility over the investment horizon as an investor that rebalances every month. Similar as in the case with unconditional returns and risk aversion larger than one, we solve the wealth equivalent of the utility loss in the case of conditional returns numerically.

4.3 Suboptimal start portfolio and transaction costs

Similar to the case with unpredictable returns we evaluate the impact of a suboptimal initial portfolio and proportional transaction costs. Table 5.6 shows the results in panel B and C. We find that the utility loss due to a suboptimal initial portfolio assuming predictable returns compares to the impact in the IID case. Not surprisingly, transaction costs increase the utility cost of rebalancing frequently. The impact of transaction costs is slightly different compared to the IID case due to the small difference in portfolio turnover. Table 5.5 shows the average annual turnover in the case of predictable returns compared to Table 5.1 for the unpredictable returns. On a single horizon ($n = 1$) we find that predictability lowers the turnover slightly. This implies that the distribution of the monthly return differences between stock and bond indices as implied by the VAR is somewhat different.

4.4 Understanding the impact of predictability

We find that predictability of asset returns influences the impact of rebalancing. It is not easy to understand these dynamics intuitively based on the results so far. In order to identify the main drivers of the differences between results based on predictable and unpredictable returns we evaluate each variable in the VAR separately, denoted by $\boldsymbol{u}_t$. We classify each variable as contributing to trending (T) or mean reversion (MR) dynamics of the equity and bond index. Because of the higher volatility, the dynamics in the returns of the equity index are dominant for the impact on rebalancing. We show that the utility gain of quarterly and yearly rebalancing in stead of monthly rebalancing is caused by multiple variables used in our VAR analysis. These together have more impact than the dividend-price ratio, of which we show (and know) that it has a mean reverting character.

In order to classify each variable we look at the sign of the regression coefficient b:

$$\boldsymbol{e}_t = a + b\boldsymbol{u}_{t-1} + \boldsymbol{v}_t, \tag{18}$$

and the correlation ρ between $\boldsymbol{v}_t$ and $\boldsymbol{\eta}_t$. $\boldsymbol{\eta}_t$ is the residual in the autoregressive model:

$$\boldsymbol{u}_t = a + b\boldsymbol{u}_{t-1} + \boldsymbol{\eta}_t. \tag{19}$$

While the correlation coefficient shows the immediate relation between surprises of the equity index and the explanatory variable, the regression coefficient shows the impact of this move on the expected return of the equity index for the following period. For example, if the correlation ρ is positive, this means that if there is a positive surprise in the explanatory variable, the excess return on the stock market is likely to be positive as well. If then the regression coefficient b is positive, the upside surprise also increases the expected return of the equity index the next period, resulting in a trend. So, if both coefficients are either positive or negative, this variable contributes to trending behaviour. If, however, both signs are different, the variable contributes to mean reversion. We do this for both the equity index and the bond index. Table 5.7 shows the results of this analysis. The dividend price ratio contributes to mean reverting behaviour of the stock market but this is dominated by the variables contributing to trending behaviour for the horizon of the rebalancing. The effect in the bond market is less clear as there are several variables contributing to mean reversion, but also several variables contributing to trends. For the impact on rebalancing, the dynamics of the equity market is more important than the dynamics of the bond market, as equity returns are the main driver of changes in portfolio weights.

5 Conclusion

For long-term passive investors, the rebalancing strategy is one of the few active decisions they have to make. For these investors, trades initiated by rebalancing are a major part of the portfolio turnover and can result in a significant amount of trading costs, lowering returns.

In this study we quantify the utility costs of periodic rebalancing of portfolios that consist of multiple risky assets, in our case equity and government bonds. Our main finding is that, although continuous rebalancing is optimal, the utility cost of less frequent rebalancing when market returns are unpredictable is limited to only basis points in terms of the initial wealth of the investor over a twenty-year period. For a simplified case we show analytically that not rebalancing leads to both drift and variation in portfolio weights. This results in an increasing utility cost of the multi-asset portfolio when the rebalancing period grows. Economically the impact is small. This cost turns into a gain if we take transaction costs into account.

Our results suggest that reducing transaction costs by less frequent rebalancing clearly outweigh optimizing the portfolio more frequently, unless transaction costs are extremely small. Using simulations, we show that these results hold for different levels of risk aversion and expected returns. Our calculations show that, when taking transaction costs into account, most of the marginal gain is captured when limiting the rebalancing frequency to once every year.

When we study a setting where asset returns are predictable, we find that a long-term investor that ignores this predictability underestimates the benefit of less frequent rebalancing. If expected market returns are time-varying there is an additional utility effect as a result of trends or mean reversion. Dependent on the time series dynamics, this effect can be positive or negative. Using a VAR model to model the predictability, and ignoring transaction costs, we estimate that the utility gain of rebalancing once every year compared to once every month is about 1.5 percent in terms of the investors' initial wealth over a twenty year period. Hence by less frequent rebalancing the investor profits from trending behavior of market returns. About half this benefit is already captured by limiting the rebalancing frequency to once every quarter instead of monthly, even when transactions costs are ignored. Time-variation of expected returns also results in different expected turnover numbers of the portfolio.

In this study we assume that our passive investor believes that market returns are IID distributed. For further research it would be interesting to look at an investor that takes predictability into account when setting the rebalancing strategy. As we show, introducing predictability in asset returns impacts the optimal portfolio. In that case, rebalancing becomes part of an active strategy and the optimal portfolio is now a function of the rebalancing period.

Tables

Table 5.1: **Portfolio turnover in the case of unpredictable returns**

This tables shows the total annualised turnover of the portfolio because of rebalancing (expressed as a fraction of the initial wealth). The standard devation shown is the standard deviation of the annualised traded volume in our simulations (number of simulations is 50000.)

	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
Average	0.40	0.23	0.16	0.11	0.10	0.07
Std.	0.13	0.08	0.06	0.04	0.04	0.03

Table 5.2: **Impact of periodic rebalancing on the expected utility in the case of unpredictable returns**

This table shows the impact of periodic rebalancing on the expected utility in the case of unpredictable returns. Panel A shows the results for an investor with log utility, ignoring transaction costs. Panel B shows the results when the initial portfolio deviates from the optimal portfolio (60% equity and 40% bonds). The assumed start portfolio is 70% equity and 30% bonds. Similar as in Panel A, transaction costs are ignored. Panel C shows the situation of Panel A only with proportional one way transaction costs of 0.1%. In each panel the rows show the absolute value of the expected utility (U-abs.), the utility cost in terms of initial wealth (SCE) and the indexed change in utility (U-index) compared to the expected utility when rebalancing continuously ($n = 1$). Between brackets are the standard errors. Number of simulations is 50000 and the investment horizon K is 240 months.

Panel A: $\gamma = 1$, no transaction costs						
	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
U-abs.	0.9683	0.9683	0.9682	0.9681	0.9679	0.9678
	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)
SCE	0.000	0.006	0.016	0.027	0.046	0.052
	(0.000)	(0.002)	(0.004)	(0.005)	(0.006)	(0.008)
Panel B: $\gamma = 1$, TC = 0.0%, different start weigths						
	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
U-abs.	0.9683	0.9682	0.9681	0.9679	0.9676	0.9675
	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)
SCE	0.000	0.007	0.021	0.041	0.071	0.083
	(0.000)	(0.003)	(0.005)	(0.007)	(0.008)	(0.010)
Panel C: $\gamma = 1$, TC = 0.1%						
	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
U-abs.	0.9639	0.9657	0.9664	0.9669	0.9669	0.9670
	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)	(0.0012)
SCE	0.000	−0.186	−0.253	−0.298	−0.299	−0.314
	(0.000)	(0.002)	(0.004)	(0.005)	(0.006)	(0.007)

Table 5.3: **Equivalent of initial wealth for different levels of risk aversion and rebalancing period in the case of unpredictable returns**

This table shows equivalent of initial wealth for different levels of risk aversion and rebalancing period. The first three colums show the impact of the rebalancing period on the expected utility for different values of risk aversion γ and constant expected returns. The last three columns show the impact of the rebalancing period on the expected utility for different expected returns, assuming constant optimal portfolios. The number of simulations is 50000 and the investment horizon is 240 months.

	constant expected return varying weights			varying expected returns constant weights		
γ	1	3	6	1	3	6
$\boldsymbol{\alpha}_e$	60%	20%	10%	60%	60%	60%
$\boldsymbol{\alpha}_b$	40%	13%	7%	40%	40%	40%
$\mathrm{E}R_e$	5.9%	5.9%	5.9%	5.9%	9.0%	13.7%
$\mathrm{E}R_b$	4.6%	4.6%	4.6%	4.6%	5.3%	6.2%
$n = 1$	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
$n = 3$	0.006%	0.007%	0.004%	0.006%	0.016%	0.047%
$n = 6$	0.016%	0.011%	0.007%	0.016%	0.039%	0.093%
$n = 12$	0.027%	0.020%	0.014%	0.027%	0.076%	0.172%
$n = 18$	0.046%	0.032%	0.022%	0.046%	0.121%	0.270%
$n = 24$	0.052%	0.047%	0.031%	0.052%	0.149%	0.309%

Table 5.4: **VAR coefficient estimates**

This table shows the regression coefficients of equation (14) and, in the last column, the R-squared for each variable. The used variables are the real short rate r_t, the stock return e_t and the 10 year treasury return b_t, the annual equity return ey_t, the nominal short rate n_t, the dividend price ratio D/P_t, the yield spread S_t and the CP factor CP_t, introduced by Cochrane and Piazzesi (2005).

	c	r_t	e_t	b_t	ey_t	n_t	D/P_t	S_t	CP_t	R^2
r_{t+1}	−0.003	0.488	−0.013	0.016	0.001	0.221	−0.001	0.000	0.000	0.409
	(0.002)	(0.039)	(0.003)	(0.006)	(0.001)	(0.071)	(0.000)	(0.000)	(0.000)	
e_{t+1}	0.071	0.670	0.053	0.162	0.006	−2.574	0.016	−0.001	0.000	0.028
	(0.029)	(0.632)	(0.047)	(0.095)	(0.013)	(1.146)	(0.007)	(0.002)	(0.001)	
b_{t+1}	−0.010	0.700	−0.090	0.086	−0.006	−0.011	−0.002	0.003	0.000	0.074
	(0.014)	(0.305)	(0.023)	(0.046)	(0.006)	(0.552)	(0.003)	(0.001)	(0.001)	
ey_{t+1}	0.060	−0.035	0.154	0.284	0.922	−2.093	0.014	0.001	−0.001	0.870
	(0.040)	(0.867)	(0.064)	(0.130)	(0.017)	(1.573)	(0.009)	(0.003)	(0.002)	
n_{t+1}	0.004	−0.006	−0.002	−0.003	0.001	0.841	0.001	0.000	0.000	0.949
	(0.000)	(0.009)	(0.001)	(0.001)	(0.000)	(0.017)	(0.000)	(0.000)	(0.000)	
D/P_{t+1}	−0.029	−0.753	−0.445	−0.158	0.004	0.872	0.993	0.001	−0.001	0.995
	(0.020)	(0.434)	(0.032)	(0.065)	(0.009)	(0.786)	(0.005)	(0.002)	(0.001)	
S_{t+1}	2.472	12.421	1.387	0.143	−0.159	−109.845	0.415	0.341	0.313	0.605
	(0.516)	(11.324)	(0.837)	(1.702)	(0.225)	(20.534)	(0.120)	(0.040)	(0.026)	
CP_{t+1}	−0.845	1.627	0.630	9.724	−0.137	68.861	−0.163	0.224	0.723	0.690
	(0.641)	(14.052)	(1.039)	(2.112)	(0.280)	(25.481)	(0.149)	(0.050)	(0.032)	

Table 5.5: **Portfolio turnover in the case of predictable returns**

This table shows the total annualised turnover of the portfolio because of rebalancing (expressed as a fraction of the initial wealth) in the case of predictable returns. The standard devation shown is the standard deviation of the annualised traded volume in our simulations (number of simulations is 50000.)

	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
Average	0.39	0.24	0.17	0.12	0.10	0.08
Std.	0.10	0.07	0.05	0.04	0.04	0.03

Table 5.6: **Impact of periodic rebalancing on the expected utility in the case of predictable returns**

This table shows the impact of periodic rebalancing on the expected utility in the case of predictable returns. Panel A shows the results for an investor with log utility, ignoring transaction costs. Panel B shows the results when the startportfolio deviates from the optimal portfolio (60% equity and 40% bonds). The assumed start portfolio is 70% equity and 30% bonds. Similar as in Panel A, transaction costs are ignored. Panel C shows the situation of Panel A only with proportional one way transaction costs of 0.1%. The rows show the absolute value of the expected utility (U-abs.) and the utility cost in terms of initial wealth (SCE) compared continuous rebalancing ($n = 1$). Between brackets are the standard errors. Number of simulations is 50000 and the investment horizon K is 240 months.

Panel A: $\gamma = 1$, no transaction costs						
	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
U-abs.	0.9521	0.9609	0.9621	0.9629	0.9629	0.9627
	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)
SCE	0.000	−0.875	−0.991	−1.077	−1.072	−1.049
	(0.000)	(0.004)	(0.007)	(0.010)	(0.013)	(0.015)
Panel B: $\gamma = 1$, TC = 0.0%, different start weigths						
	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
U-abs.	0.9522	0.9609	0.9621	0.9628	0.9627	0.9624
	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)
SCE	0.000	−0.865	−0.984	−1.056	−1.045	−1.012
	(0.000)	(0.005)	(0.009)	(0.013)	(0.016)	(0.018)
Panel C: $\gamma = 1$, TC = 0.1%						
	$n = 1$	$n = 3$	$n = 6$	$n = 12$	$n = 18$	$n = 24$
U-abs.	0.9477	0.9582	0.9602	0.9616	0.9618	0.9618
	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)
SCE	0.000	−1.047	−1.243	−1.386	−1.402	−1.402
	(0.000)	(0.004)	(0.007)	(0.010)	(0.013)	(0.015)

Table 5.7: **Understanding time series dynamics**

This tables shows the regression coefficient (first row) and the R-squared (second row) for every explanatory variable separately (equation 18). The third row shows the correlation of the residuals $\boldsymbol{v}_t$ and $\boldsymbol{\eta}_t$). The last row shows the classification of each variable. Panel A shows the above for the equity index; Panel B for the government bond index. The used variables are the real short rate r_t, the stock return e_t and the 10 year treasury return b_t, the annual equity return ey_t, the nominal short rate n_t, the dividend price ratio D/P_t, the yield spread S_t and the CP factor CP_t.

Panel A: Equity index							
	e_t	b_t	r_t	n_t	D/P_t	S_t	CP_t
b	0.0720	0.2058	0.3721	−0.6481	0.0046	0.0004	0.0000
ρ	1.0000	0.0899	0.0419	−0.0034	−0.6777	0.1061	0.0527
R^2	0.52%	1.03%	0.09%	0.17%	0.20%	0.01%	0.00%
T/MR	T	T	T	T	MR	T	T
Panel B: Bond index							
	e_t	b_t	r_t	n_t	D/P_t	S_t	CP_t
b	−0.0815	0.0684	0.7046	−0.1946	−0.0003	0.0023	0.0011
ρ	0.1099	1.0000	0.1304	−0.0263	0.0367	−0.3239	−0.1340
R^2	2.72%	0.47%	1.29%	0.06%	0.00%	1.73%	0.84%
T/MR	MR	T	T	T	MR	MR	MR

Figures

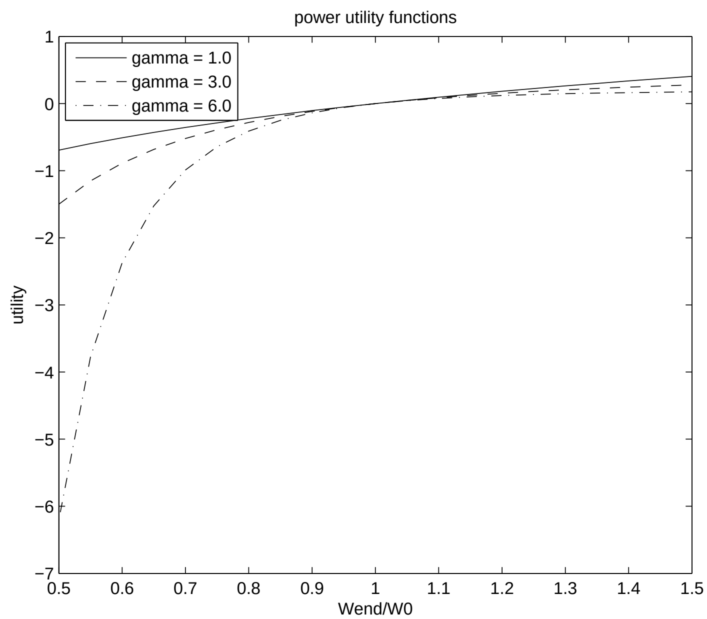

Figure 5.1: **Power utility function for different values of risk aversion**

The figure shows the power utility function for different values of risk aversion γ. Since we assume a constant relative risk aversion (CRRA) for our investor, it is the ratio between the final wealth (W_{end}) and initial wealth (W_0) that matters for the utility.

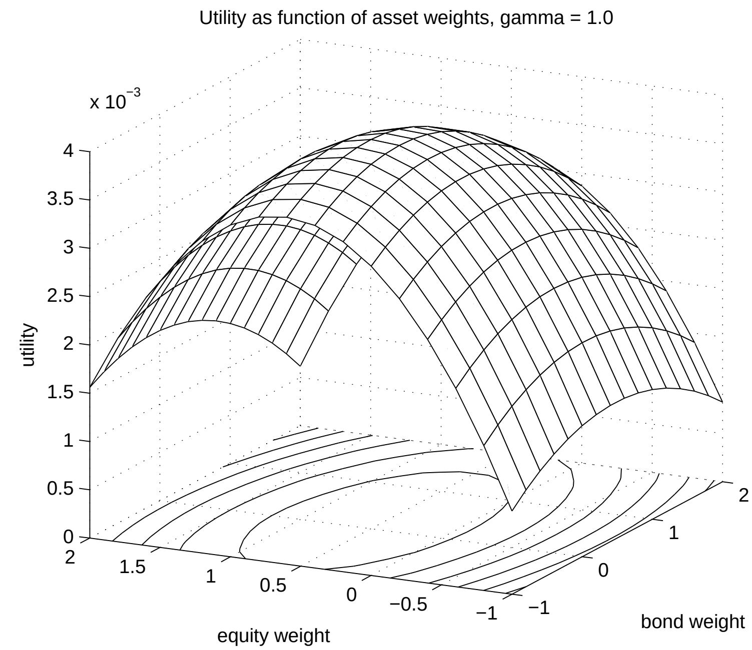

(a)

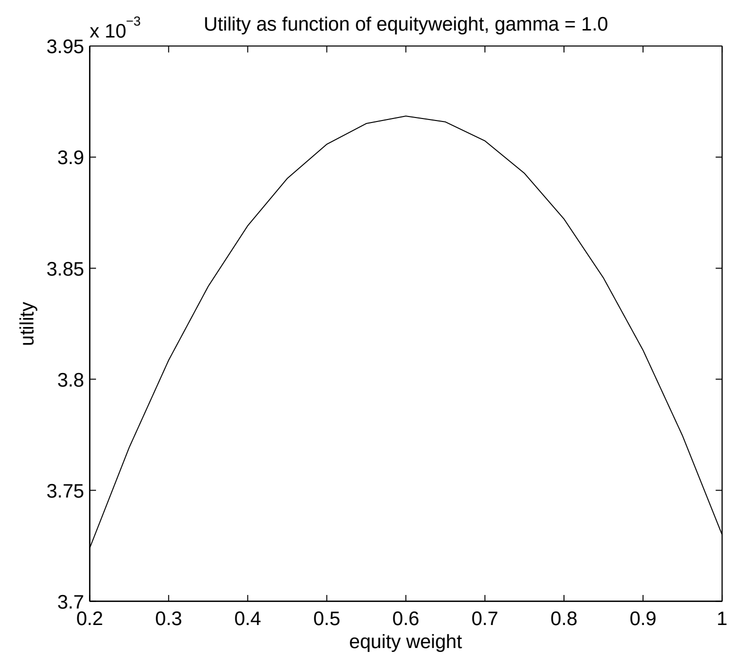

(b)

Figure 5.2: **Utility as function of portfolio weight of risky assets**

These figures shows utility as function of portfolio weight of risky assets. Figure 5.2a shows utility as function of both equity and bond weight. Figure 5.2b shows utility as function of equity weight, assuming an fully invested, unlevered portfolio of bonds and equity.

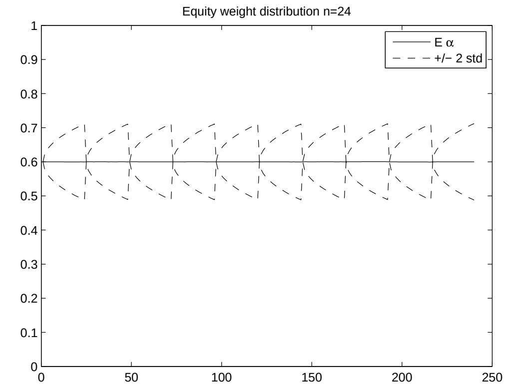

(a)

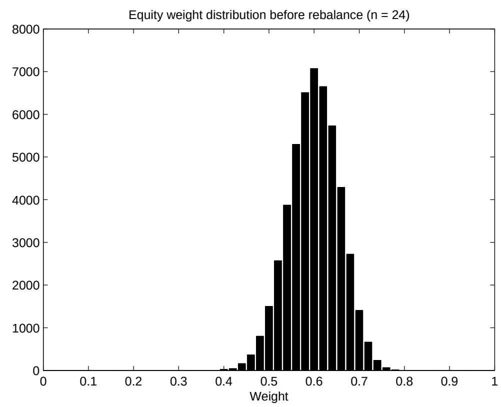

(b)

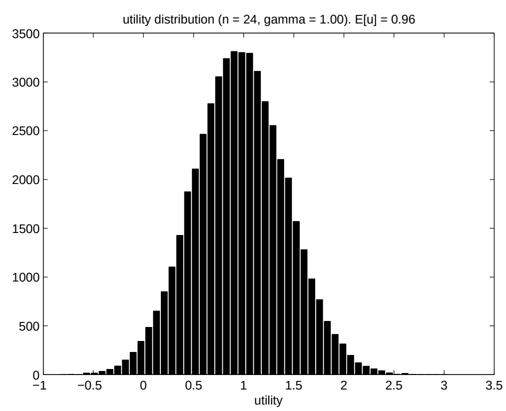

(c)

Figure 5.3: **Simulation results when rebalancing every 24 months**

These figures show the simulation results when rebalancing every 24 months. Figure 5.3a shows the distribution and mean of the portfolio weight of equity over times. Figure 5.3b shows the resulting weight distribution before rebalancing, and Figure 5.3c shows the resulting utility distribution assuming log utility.

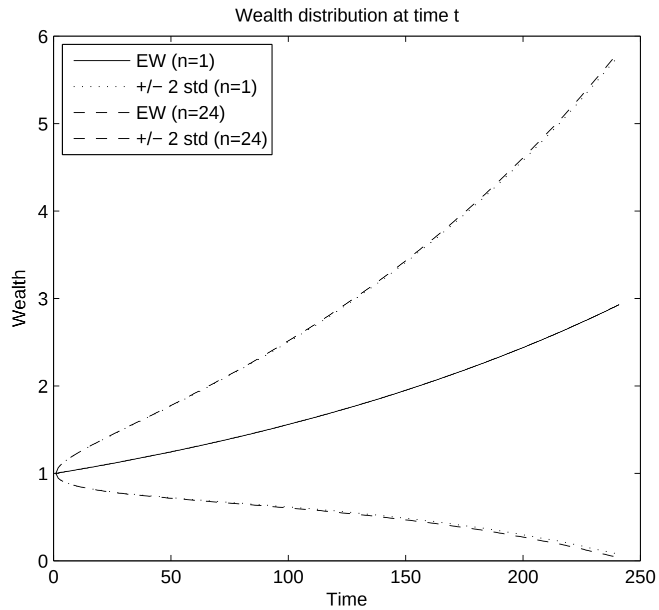

(a)

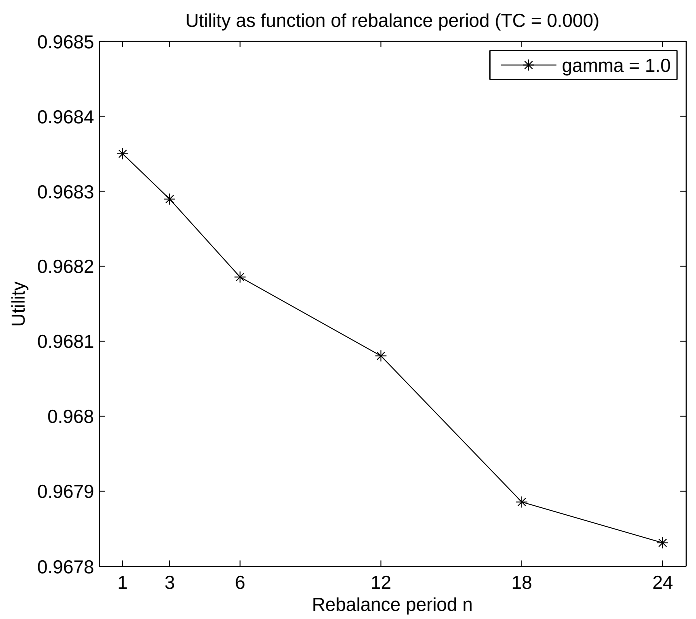

(b)

Figure 5.4: **Impact of the rebalance period on the portfolio returns and utility**

This figures shows the impact of the rebalance period on the portfolio returns and utility. Figure 5.4a shows the expected wealth and confidence interval for both continuous rebalancing ($n = 1$) and rebalancing every 24 months ($n = 24$). In both cases the horizon K is 240 months. Figure 5.4b shows the impact of the rebalance period n on the expected utility.

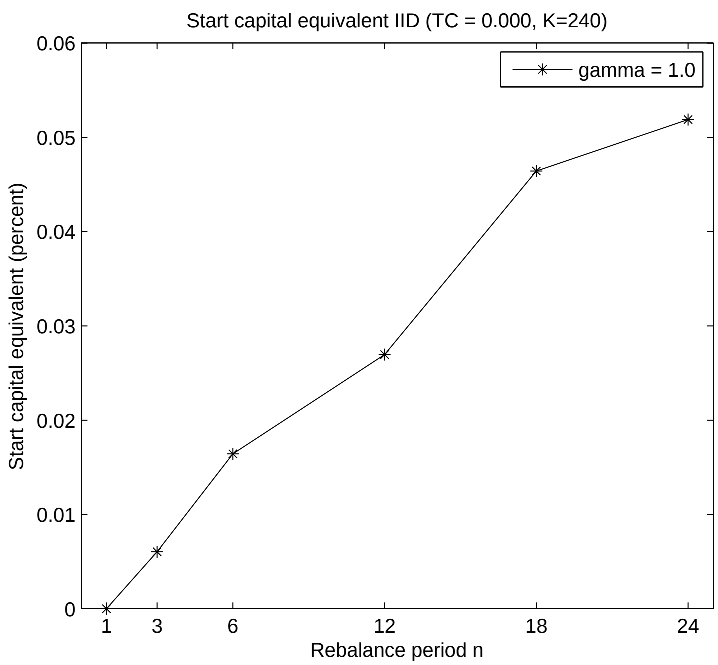

Figure 5.5: **Initial wealth equivalent of the utility loss**

The figure shows the initial wealth equivalent of the utility loss as a result of less frequent rebalancing.

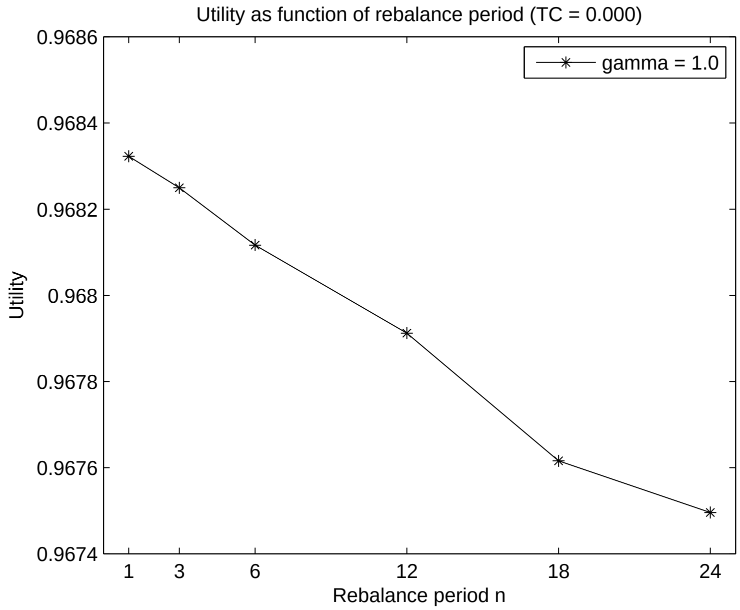

(a)

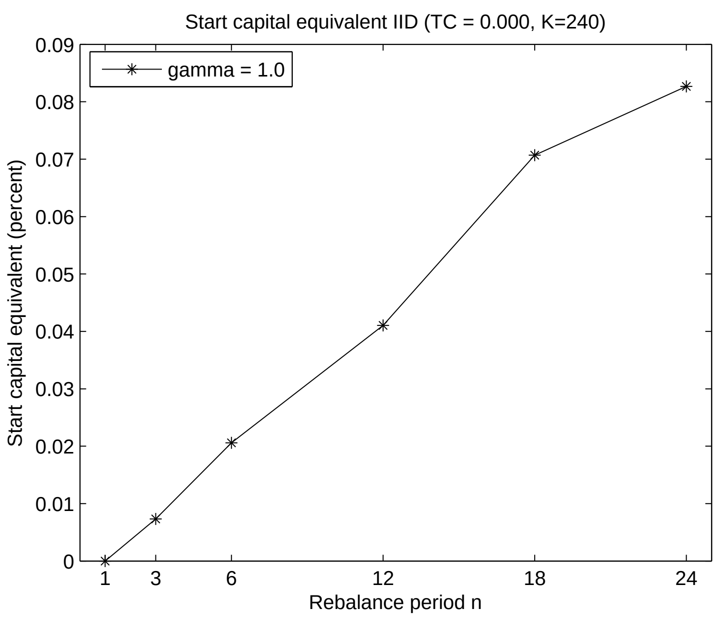

(b)

Figure 5.6: **Impact of the rebalance period on the utility assuming a suboptimal portfolio at start**

This figure shows the impact of the rebalance period on the utility assuming a suboptimal portfolio at start. Figure 5.6a shows the impact of rebalancing on utility assuming a suboptimal portfolio at start. Figure 5.6b shows the impact of the rebalance period as initial wealth equivalent assuming a suboptimal portfolio at start.

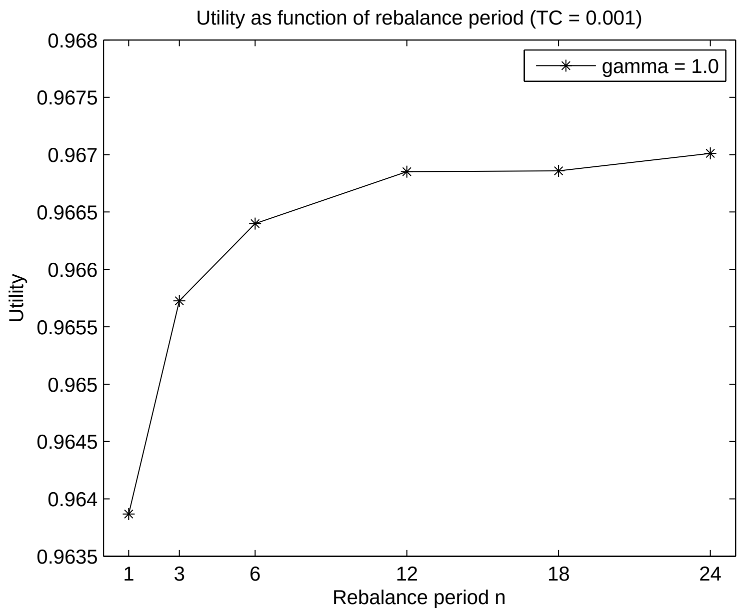

(a)

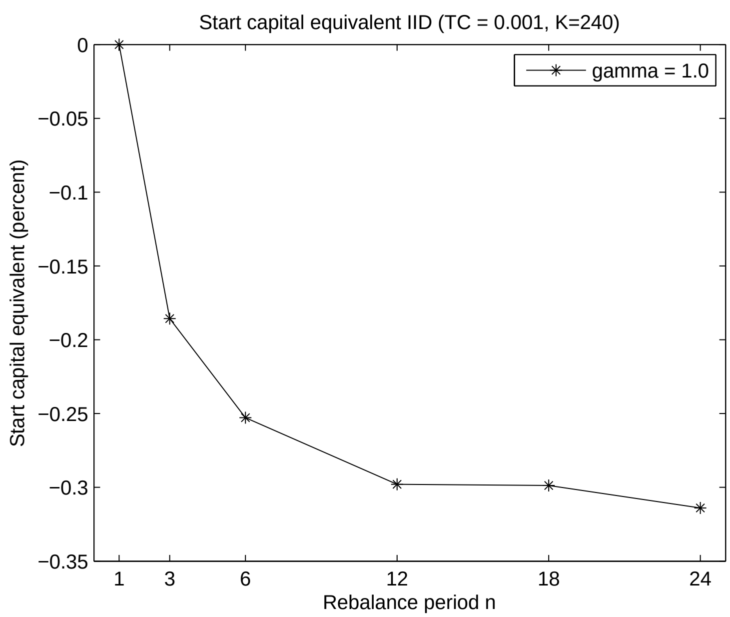

(b)

Figure 5.7: **Impact of the rebalance period on the portfolio utility**

This figure shows the impact of the rebalance period on the portfolio utility. Figure 5.4a shows the impact of rebalancing on utility, with proportional trading costs (0.1% of traded volume). Figure 5.4b shows the impact of the rebalance period as initial wealth equivalent with proportional trading costs (0.1% of traded volume).

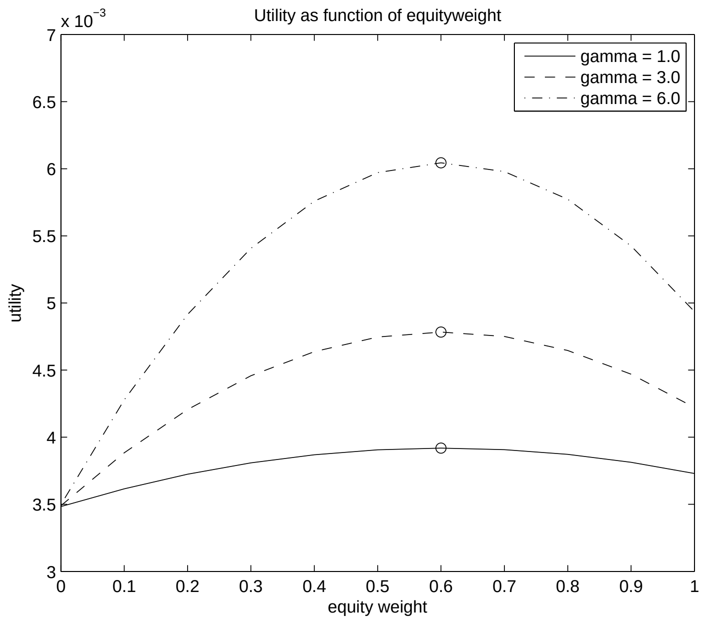

(a)

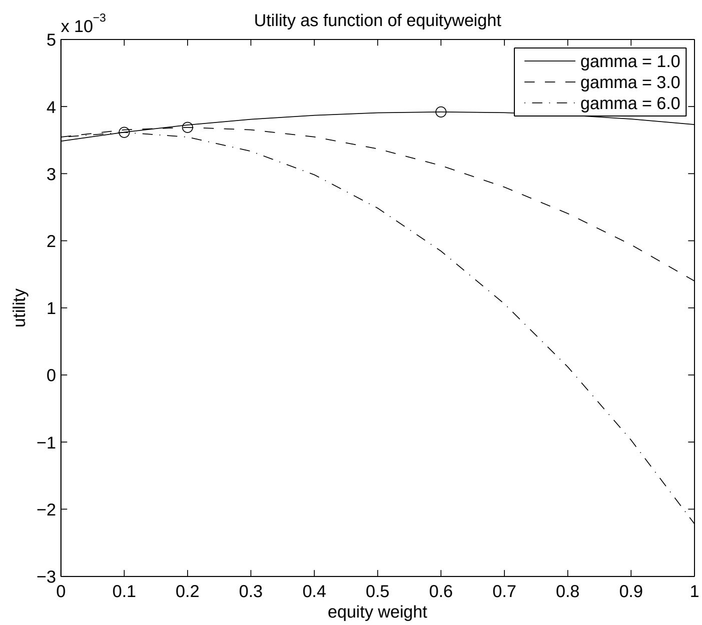

(b)

Figure 5.8: **Impact of risk aversion on utility curves**

This figure shows the impact of the level of risk aversion on the utility curves as function of the equity weight. Figure 5.8a shows the utility as function of equity weight and different risk premiums. Figure 5.8b shows the utility as function of equity weight and different optimal portfolio weights.

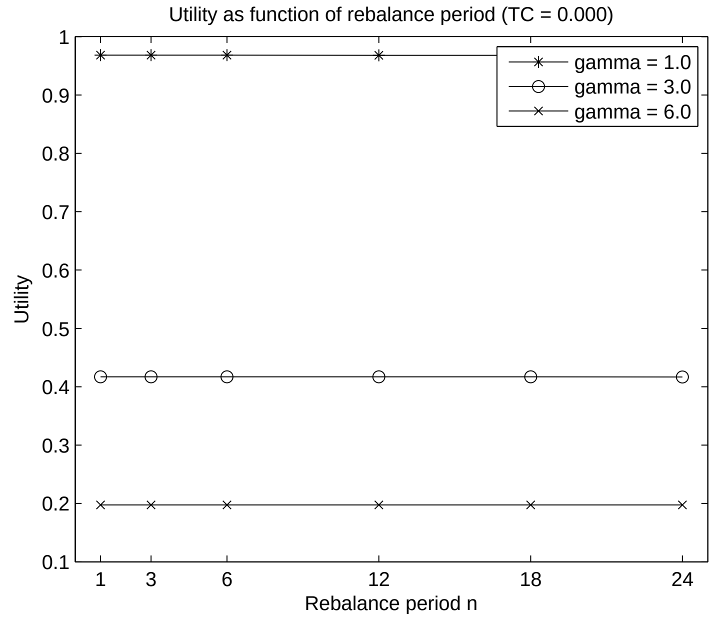

(a)

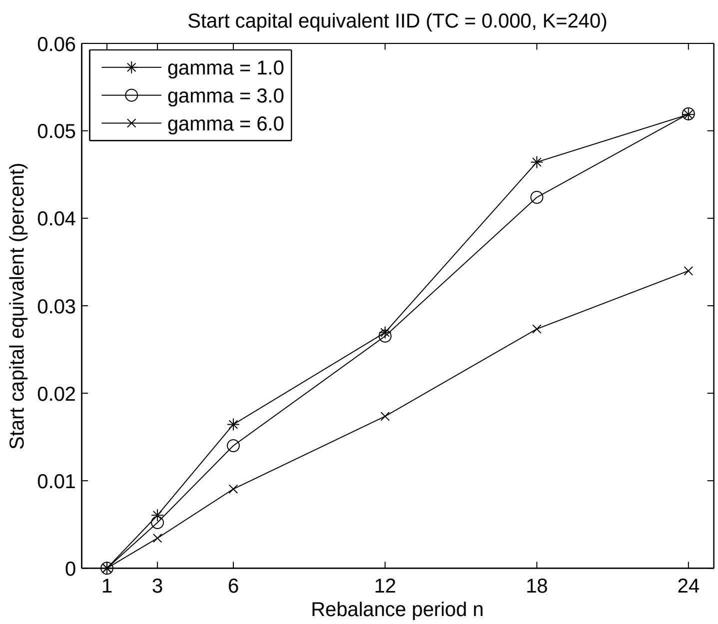

(b)

Figure 5.9: **The impact of the level of risk aversion and portfolio weights**

This figure shows the impact of the level of risk aversion and portfolio weights on expected utility and initial wealth equivalent. Figure 5.9a shows the impact of rebalancing on the expected utility for different levels of risk aversion. The optimal allocation varies, the expected risk premium is constant for each level. Figure 5.9b shows the impact of rebalancing in terms of initial wealth equivalent for different levels of risk aversion. The optimal allocation varies, the expected risk premium is constant for each level.

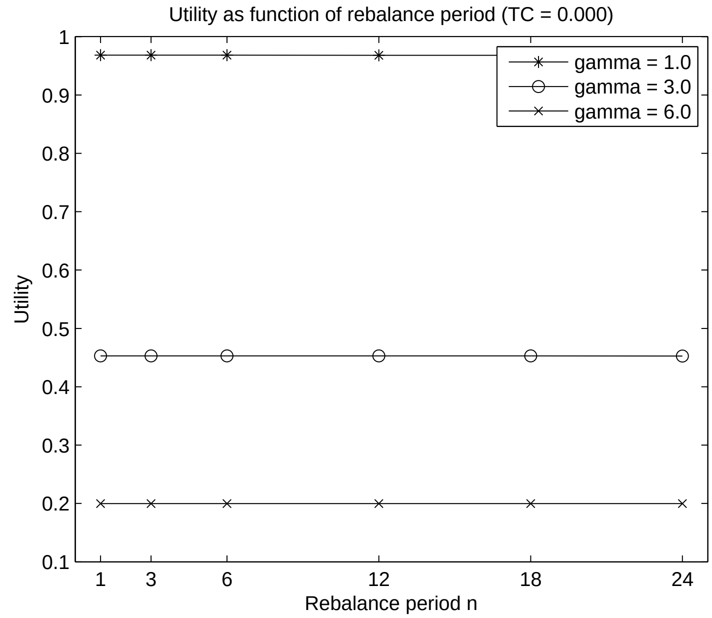

(a)

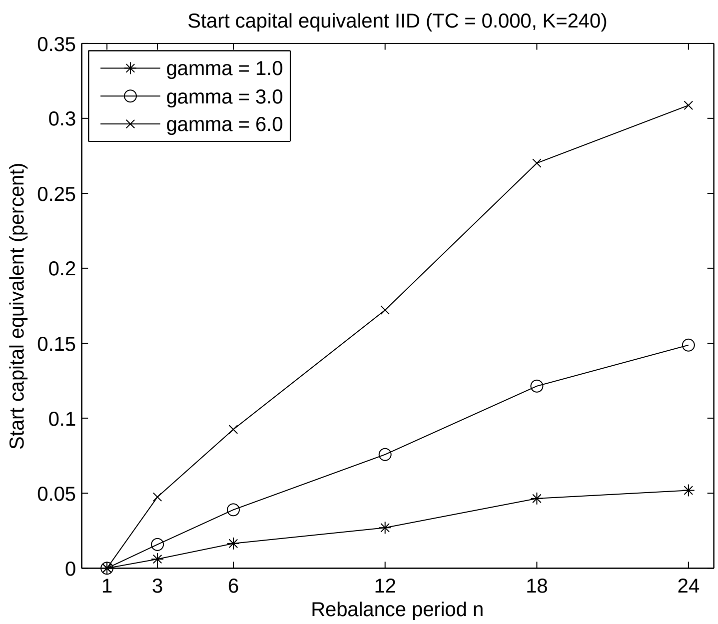

(b)

Figure 5.10: **Impact of the level of risk aversion and expected risk premium**

This figure shows the impact of the level of risk aversion and expected risk premium on expected utility and initial wealth equivalent. Figure 5.10a shows the impact of rebalancing on the expected utility for different levels of risk aversion. The optimal allocation is constant, the expected risk premium is different for each level. Figure 5.10b shows the impact of rebalancing on the initial wealth equivalent for different levels of risk aversion. The optimal allocation is constant, the expected risk premium is different for each level.

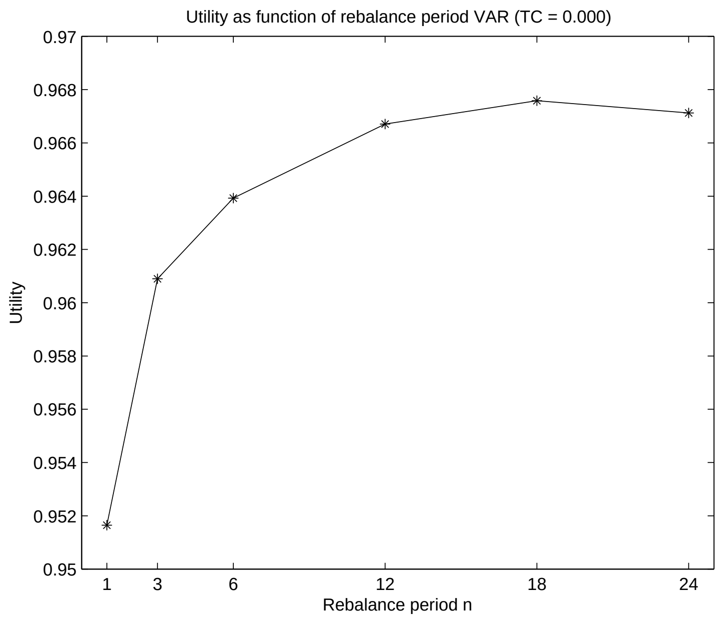

(a)

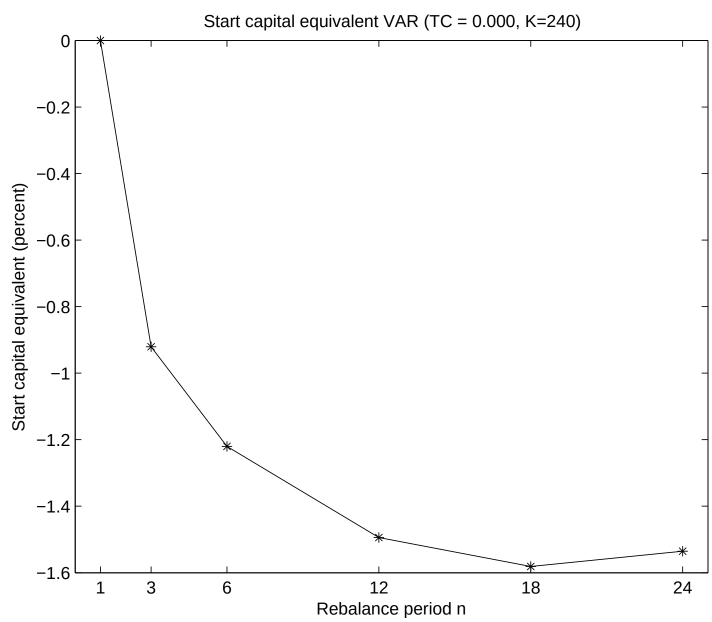

(b)

Figure 5.11: **Impact of the rebalance period on the portfolio in the case of predictable returns**

This figure shows the impact of the rebalance period on the portfolio returns and utility, when returns are predictable. Figure 5.11a shows the impact of the rebalance period n on the expected utility. Figure 5.11b shows the initial wealth equivalent of the utility loss by less rebalancing.

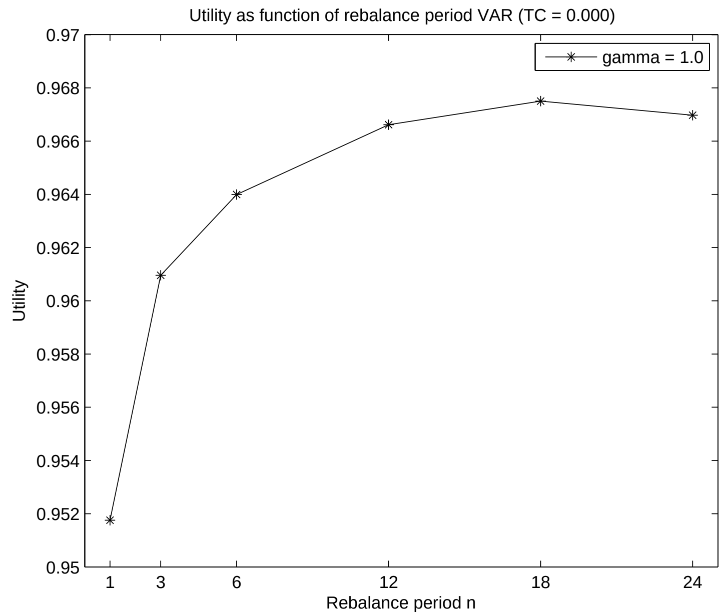

(a)

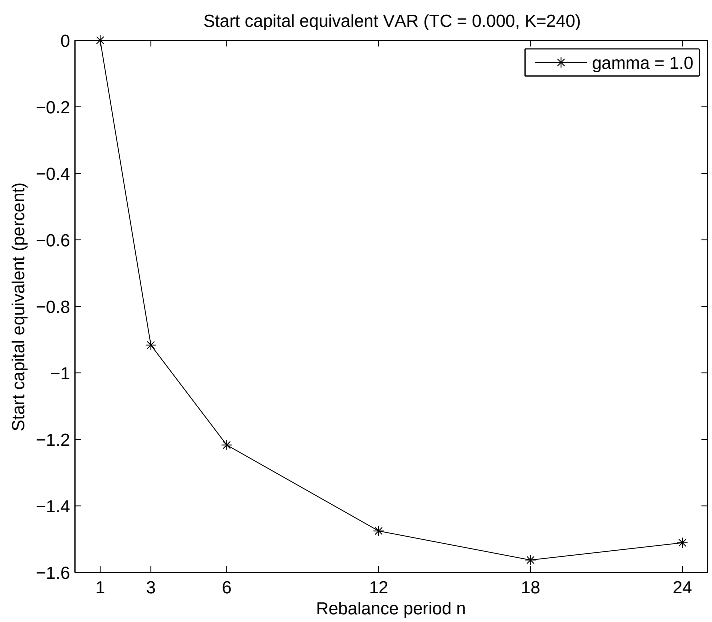

(b)

Figure 5.12: **Impact of the rebalance period in the case of predictable returns and a suboptimal portfolio at start**

This figure shows the impact of the rebalance period on the utility assuming a suboptimal portfolio (70% equity) at start. Figure 5.12a shows the impact of rebalancing on utility, assuming a suboptimal portfolio at start. Figure 5.12b shows the initial wealth equivalent as function of the rebalance period, assuming a suboptimal portfolio at start.

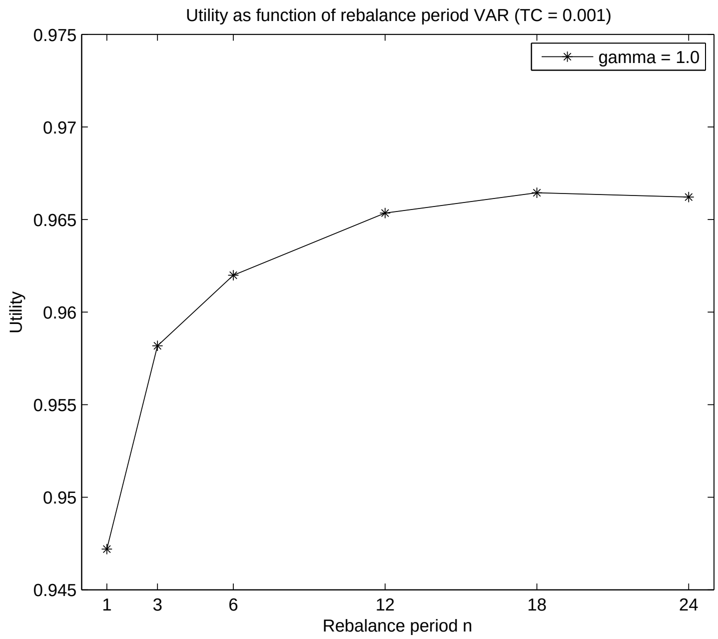

(a)

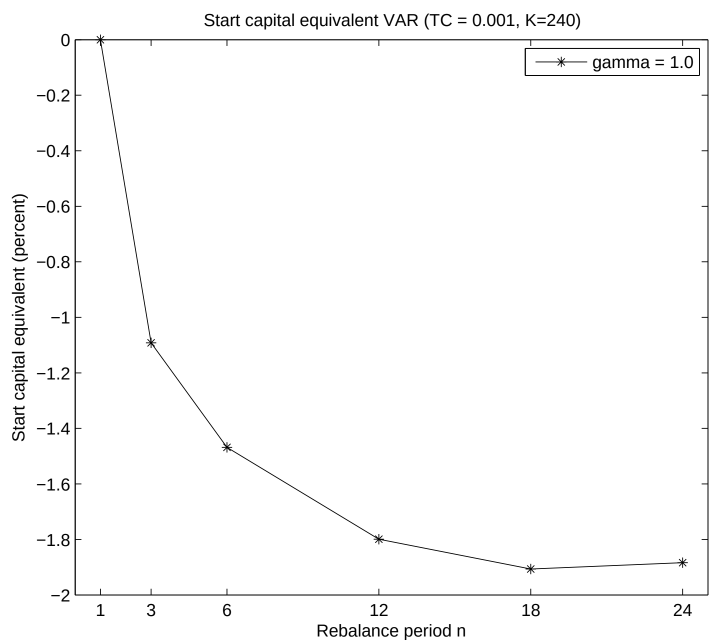

(b)

Figure 5.13: **Impact of the rebalance period in the case of predictable returns and transaction costs**

This figure shows the impact of the rebalance period on the utility and initial wealth equivalent in the case of transaction costs. Figure 5.13a shows the impact of rebalancing on utility, with proportional trading costs. Figure 5.13b shows the initial wealth equivalent as function of rebalance period, with proportional trading costs.

References

Ang, Andrew. *Asset management: A systematic approach to factor investing.* Oxford University Press, 2014.

Balduzzi, Pierluigi and Anthony W Lynch. Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics*, 52(1):47–78, 1999.

Bikker, Jacob A, Laura Spierdijk, and Pieter Jelle Van Der Sluis. Market impact costs of institutional equity trades. *Journal of International Money and Finance*, 26(6):974–1000, 2007.

Bouchey, Paul, Vassilii Nemtchinov, Alex Paulsen, and David M Stein. Volatility harvesting: Why does diversifying and rebalancing create portfolio growth? *The Journal of Wealth Management*, 15(2):26, 2012.

Campbell, John Y and Luis M Viceira. *Strategic asset allocation: portfolio choice for long-term investors.* Oxford University Press, USA, 2002.

Campbell, John Y, Yeung Lewis Chan, and Luis M Viceira. A multivariate model of strategic asset allocation. *Journal of financial economics*, 67(1):41–80, 2003.

Chambers, David, Elroy Dimson, and Antti Ilmanen. The norway model. *The Journal of Portfolio Management*, 38(2):67–81, 2012.

Cochrane, John H. *Asset pricing.* {Princeton University Press}, 2001.

Cochrane, John H and Monika Piazzesi. Bond risk premia. *American Economic Review*, 95(1): 138–160, 2005.

Constantinides, George M. Capital market equilibrium with transaction costs. *The Journal of Political Economy*, pages 842–862, 1986.

Erb, Claude B and Campbell R Harvey. The strategic and tactical value of commodity futures. *Financial Analysts Journal*, 62(2):69–97, 2006.

Liu, Hong. Optimal consumption and investment with transaction costs and multiple risky assets. *The Journal of Finance*, 59(1):289–338, 2004.

Merton, Robert C. Lifetime portfolio selection under uncertainty: The continuous-time case. *The review of Economics and Statistics*, pages 247–257, 1969.

Merton, Robert C. Optimum consumption and portfolio rules in a continuous-time model. *Journal of economic theory*, 3(4):373–413, 1971.

Mulvey, John M and Koray D Simsek. Rebalancing strategies for long-term investors. *Applied Optimization*, 74:15–31, 2002.

Perold, Andre F and William F Sharpe. Dynamic strategies for asset allocation. *Financial Analysts Journal*, 51(1):149–160, 1995.

Samuelson, Paul A. Lifetime portfolio selection by dynamic stochastic programming. *The review of economics and statistics*, pages 239–246, 1969.

Sharpe, William F. Adaptive asset allocation policies. *Financial Analysts Journal*, 66(3):45–59, 2010.