Erasmus University Rotterdam
Erasmus School of Economics
Bachelor Thesis Econometrics and Operations Research

Simulation-Based Dynamic Programming: A Study on Portfolio Performance

Daniel Moberg (580512)

Supervisor:	Kole, HJWG
Second assessor:	Vladimirov, E
Date final version:	1st July 2024

Abstract

This paper examines the portfolio performance of the BGSS method, proposed by Brandt et al. (2005) to evaluate its profitability for investors and the potential benefits of incorporating multiple state variables. The method is performed on three different state variables: the log dividend price ratio, the default yield spread, and the return on long-term government bonds as well as combinations of these. Individually, the BGSS model with the log dividend price ratio as the state variable is consistently able to outperform the $1/N$ portfolio for an investment horizon of 8 quarters or more. For risk-averse investors, the BGSS method with the default yield spread or long-term return as the state variable can be appealing due to its significantly lower return variance and its more stable Sharpe Ratios over time compared to the $1/N$ model. Furthermore, we find that combinations including a well-performing state variable, particularly the log dividend price ratio, increase results slightly but not significantly. These findings suggest that the BGSS method can be beneficial, but careful selection of state variables is crucial for optimal performance.

1 Introduction

Dynamic programming methods have been utilized in portfolio optimization due to their ability to take future expectations into account. Dynamic programming methods can be beneficial for multi-period investing: investing from time t until T with the option to rebalance the portfolio at each point in between time t and $T-1$. Institutions such as pension funds can benefit from employing this approach, as its primary goal is to provide retirement benefits to its customers over many years or even decades. However, dynamic programming is also relevant for short-term investors aiming for optimal results within a shorter investment period of a minimum of two time periods. This paper evaluates the portfolio performance of such a dynamic programming method, namely the BGSS method. By implementing this method in an investment scenario where an investor allocates his wealth between one risky asset and the risk-free rate, the performance is assessed using various performance measures.

The main goal of this paper is to examine the out-of-sample portfolio performance of the BGSS method in a real-world investment scenario compared to the $1/N$ strategy. Additionally, we determine whether incorporating different state variables simultaneously can improve the results. This study aims to contribute to the understanding of the BGSS method in portfolio optimization by evaluating its practical advantages.

In Van Binsbergen & Brandt (2007), inspired by Brandt et al. (2005) this dynamic programming method is presented. The BGSS method differs from a traditional approach, discretizing the state space. At each point of several simulated paths, a regression of the utilities on the simulated state variables is performed. These fitted regression values provide the expected utility for each path and are therefore used to calculate the optimal weights. Van Binsbergen & Brandt (2007) compares the results of either iterating on the optimal portfolio weights (portfolio iteration) or iterating on the value function (value iteration). To evaluate which method provides more accurate results, they compare the obtained weights at time 0 with those obtained with the discrete state space method, which serves as their benchmark. This comparison raises the question of the BGSS method's usefulness if the discrete state space method serves as the benchmark. Specifically, what advantages does the BGSS method offer over the discrete state space method? Additionally, they only examine the weight at time 0 and do not address the real out-of-sample portfolio performance. In this paper, we will address this question by applying the BGSS method on a whole investment period (obtaining weights for every quarter) and evaluating the out-of-sample performance.

A significant advantage of the BGSS method compared to discretizing the state space is that it is less prone to the curse of dimensionality. Let's assume instead of using only one state variable, the number of state variables is k. In the discrete state space method, n number of grid points is defined for a state variable. However, if the number of state variables changes to k, the number of grid points increases to n^k. This exponential growth can result in either computational inefficiency when the number of grid points is not reduced or in a more imprecision method when deciding to reduce the number of grid points. Moreover, interpolation becomes tricky when the number of state variables grows. In contrast, the BGSS method only simulates a set of paths of the returns and the state variables without defining grid points or using interpolation, therefore avoiding these issues.

In this paper, we will also evaluate to what extent this advantage is practically significant. By using more state variables individually and simultaneously, we can investigate if the difference in portfolio performance among the state variables and combinations is significant. This allows for the investigation if using multiple or different state variables can provide practical advantages for investors. In theory, the advantage of using more state variables seems an improvement since it could improve the predictive power and the fitness of the model. However, multiple studies (such as Berry & Feldman (1985)) have suggested that using multiple regressions can also cause problems, such as multicollinearity.

The BGSS method is assessed using three different state variables: the log dividend price ratio (LDP), the return on long-term government bonds (LTR), and the default yield spread (DYS). Additionally, to evaluate the potential advantage of using multiple state variables, we apply the BGSS method with different combinations of these variables. The models are performed on different levels of risk aversion ($\gamma = 5, 10, 15, 20$) and different investment horizons ($T = 2, 4, 8, 12, 20, 40$). The results are compared to a $1/N$ benchmark model and a one-period BGSS benchmark model using different performance measures, namely the returns and variance of the portfolio, the Sharpe ratio, and the Certainty Equivalent.

Individually, the BGSS model with LDP as the state variable consistently outperforms the $1/N$ model for an investment horizon of 8 quarters or longer. For these investment horizons, the model contains at least two levels of risk aversion which provide a higher Sharpe ratio. For example, for $T = 40$ with a level of risk aversion of $\gamma = 20$, the Sharpe ratio is 0.91 compared to 0.75 for the $1/N$ model. Moreover, it consistently provides a higher Certainty Equivalent, indicating that the return an investor considers equal is higher than for the $1/N$ model. The log dividend price ratio performs significantly better as a state variable than the default yield spread and long-term return since it provides a higher Sharpe ratio and higher Certainty Equivalent for every level of risk aversion and every investment horizon of 4 quarters or higher. The obtained weights are significantly higher when the log dividend price ratio is used as the state variable for the periods before 2020 when the S&P 500 seems to perform well since the high weights result in relatively high Sharpe ratios. The result of higher simulated returns with the log dividend price ratio can partially explain this since this increases the chances of a higher expected utility after the regression. This higher expected utility results in a higher optimal weight. The higher weights result in higher mean returns which sufficiently compensate for higher obtained variance, thereby leading to good Sharpe ratios.

The BGSS models with LTR or DYS as the state variable are not able to outperform the $1/N$ consistently. In most cases, these models obtain a worse Sharpe Ratio than the $1/N$ model, except for a few expectations. Although the $1/N$ model obtains higher Sharpe ratios for most investment periods/levels of risk aversion, the BGSS models with DYS or LTR as the state variable can be appealing to more risk-averse investors due to the significantly lower portfolio variance and more stable Sharpe ratios over time compared to the $1/N$ model. Additionally, for all BGSS models, the overall performance is the best for a lower level of risk aversion (except for an investment horizon of 40 quarters), suggesting potential further improvements with lower risk aversion levels.

When using more state variables simultaneously the results increase slightly for the combinations including the LDP. However, these improvements are not significant. For example, with an investment horizon of 20 quarters and a level of risk aversion of $\gamma = 10$, the individual LDP model obtains a Sharpe ratio of 0.81 and the combinations including the LDP obtain a Sharpe Ratio of 0.86 (LDP/DYS), 0.80 (LDP/LTR) and 0.85 (LDP/LTR/DYS). The last result indicates that the combination with all three state variables does not improve the outcomes compared to the combinations with two state variables, which generally is the case. The combination with LTR and DYS as the state variables performs significantly worse, which is not surprising when looking at the individual results. Therefore the conclusion can be drawn that using combinations of state variables can be beneficial but the state variables must perform well individually to get good results. So, careful selection of state variables is crucial for obtaining optimal performance. Besides, adding state variables that underperform (LTR and DYS) compared to another state variable (LDP) can improve results slightly but not significantly.

This paper is organized as follows: In Section 2 a literature overview is given. Section 3 presents the data. Section 4 explains the BGSS method, the construction of portfolios, and the performance measures. Furthermore, in Section 5 the results are discussed and the conclusions are drawn in Section 6.

2 Literature review

Dynamic portfolio choice is a well-known approach in the literature. Dynamic portfolio choice is a strategy for managing investment portfolios that involves intermediately adjusting asset allocations over time. A popular dynamic programming approach involves discretizing the state space. Once the space is discretized the value function can be approximated using different techniques. Receptively, Brandt (1999), Balduzzi & Lynch (1999), and Barberis (2000) propose nonparametric regressions, quadrature integration, and simulations as techniques to approximate this value function. In each step of the backward recursion used in dynamic programming, the optimization problem is then solved by maximizing the expectation of this approximated value function derived in the previous period.

Brandt et al. (2005) propose a different approach where at every point in time regressions of simulated returns on simulated state variables are performed. Van Binsbergen & Brandt (2007) solve the question of whether it is optimal to iterate on the value function or directly on the portfolio weights. They find that iterating on optimized portfolio weights leads to less biased results than iterating on the value function when incorporating short-sale constraints.

Garlappi & Skoulakis (2009) argue that the poor performance of the value function iteration is due to the fact of the high nonlinearity of the value function. They suggest using the Certainty Equivalent instead of the value function to do the value iteration. With this approach, they obtained very accurate results.

Cong & Oosterlee (2017) extend the paper of BGGS by using the SGBM approach, proposed Jain & Oosterlee (2015), to calculate the conditional expectations instead of using the standard regression model used in the BGSS method. Besides, they use the Taylor expansion introduced by Garlappi & Skoulakis (2009). By substituting these methods in the BGGS method they find superior results in comparison to other simulation-based approaches.

Simulation-based dynamic programming methods can be used for asset allocation. However, outperforming the simple $1/N$ portfolio where all weights are equally assigned to every asset in the portfolio, has proven to be difficult. DeMiguel et al. (2009) presents 14 models, which all do not consistently beat the $1/N$ portfolio. This shows that a relatively simple portfolio can be highly successful. In this paper, the out-of-sample performance of the BGSS method will be examined and compared to the $1/N$ portfolio.

The aforementioned papers show that since the introduction of this simulation-based dynamic programming approach by Brandt et al. (2005), there has been significant further research on this topic. By evaluating how the method performs for our specific scenario/dataset, we aim to establish a foundation for using the BGSS method. If the findings indicate that the method provides good out-of-sample performance, it will be motivating to explore whether incorporating more recent advancements from the literature can further enhance the results.

3 Data

To evaluate the portfolio performance of the BGSS method the quarterly value-weighted returns of the S&P 500 from the period January 1994-December 2023 are used. This data is collected from the Center for Research in Security Prices (CRSP)[1]. Additionally, for the same period the quarterly data for the state variables and the risk-free rate is used. These state variables are inspired by Welch & Goyal (2008)) and are all available on Goyal's website[2]. The risk-free rate is also available in that dataset. The estimation period is 80 quarters (20 years) and the investment period depends on the time horizon (T), which varies from 2,4,8,12,20 and 40. To ensure the research is based on the most recent data, all final investment periods will end in 2023Q4, utilizing a rolling window estimation approach. For instance, when $T = 2$, the final investment period spans 2023Q3 to 2023Q4, thus the estimation sample is from 2004Q3 to 2023Q2. Moreover, when $T = 40$, the investment period extends from 2014Q1 to 2023Q4, making the estimation window 1994Q1 to 2013Q4.

[1] https://wrds-www.wharton.upenn.edu/pages/get-data/center-research-security-prices-crsp/

[2] https://sites.google.com/view/agoyal145

The different investment horizons are based on the methodology outlined in Van Binsbergen & Brandt (2007). Their estimation window for the VAR model is 40 quarters (10 years). However, we deviate from this by increasing the estimation window to 20 years. A longer estimation window can enhance the robustness and reliability of the VAR model estimates.

The state variables are the predictor variables used in Welch & Goyal (2008) and many more studies from which we know they have some level of return predictability. Besides, the state variables should have a negative concurrent correlation with the returns. If there is a negative shock to returns at time t, a positive shock of the state variable is expected. Consequently, the state variable at time t becomes larger, positively affecting future returns. This increase can compensate for any potential negative shock, indicating that these state variables have predictive power for future returns. The three state variables used in this paper are:

- 1: Log Dividend Price ratio (LDP): logarithm of a 12-month moving sum of dividends paid on the S&P 500 index divided by the current index. Since we divide by the index of the $S\&P500$, the log dividend price ratio increases if the index decreases. Therefore we expect negative concurrent correlation between the log dividend price ratio and returns.

- 2: Default Yield Spread (DYS): the difference between Moody's BAA- and AAA-rated corporate bond yields. Negative concurrent correlation is expected since a higher default yield spread reflects higher risk aversion which leads to investors moving their funds away from stocks which can lead to a decline in stock prices

- 3: Long-Term Return (LTR): return on long-term government bonds. Negative concurrent correlation is expected because typically when bond returns are high investors are moving their funds from stocks to bonds which leads to higher bond prices and lower stock prices.

The summary statistics of the returns and the state variables are shown in Table 1. The table shows a positive mean log excess return of 0.019, which is 0.078 or 7.8% annualized. This implies that if the simulations generally produce positive log excess returns, they effectively represent the real scenario. If this is the case, the models should provide significant weights to the risky asset since this results in on average positive excess return based on the positive mean. A histogram of the returns is shown in Figure 3. Based on the histogram and the JB-statistic in Table 1, we reject the null hypothesis that the returns are normally distributed. The returns are left-skewed since the skewness is negative. In other words, there are more extreme negative than extreme positive returns.

Table 2 shows the correlations between the variables. As expected, all state variables have a negative correlation with the log excess return. For example, a high log dividend price ratio typically means that stocks are undervalued and therefore the expectation is that stocks will increase. An interesting result is that the log dividend price ratio has the lowest negative correlation with the log excess return, although it has the highest negative correlation between error terms with the log excess return.

Table 1: Summary statistics quarterly log excess returns and state variables

Properties/State Variables	Log excess returns	LDP	DYS	LTR
Mean	0.019	-3.987	0.010	0.016
St Dev	0.084	0.222	0.04	0.057
Min	-0.251	-4.479	0.06	-0.116
Max	0.187	-3.342	0.034	0.211
Excess Kurtosis	1.283	3.193	19.499	4.469
Skewness	-1.028	-0.065	3.322	0.753
JB-statistic	19.570	0.269	1581.000	22.328

This table shows the summary statistics of the quarterly returns of the S&P 500 and the state variables: log dividend price ratio (LDP), default yield spread (DYS), and the return on long-term government bonds (LTR). The data sample starts from 1994Q1 and ends at December 2023Q4.

Table 2: Correlation matrix

	Returns	LDP	DYS	LTR
Returns	1.00	-0.14	-0.32	-0.40
LDP	-0.14	1.00	0.10	0.10
DYS	-0.32	0.39	1.00	0.15
LTR	-0.40	0.10	0.15	1.00

This table shows the correlation matrix with the correlations between the quarterly log excess returns on the S&P 500 and the state variables: log dividend price ratio (LDP), default yield spread (DYS), and the return on long-term government bonds (LTR). The data sample starts from 1994Q1 and ends at December 2023Q4.

4 Methodology

4.1 Portfolio Problem

The models presented in this paper are based on the following portfolio choice problem inspired by Van Binsbergen & Brandt (2007), where an investor wants to maximize its utility:

$$V_t(W_t, Z_t) = \max_{\{x_s\}_{s=t}^{T-1}} E_t[u(W_T)], \tag{1}$$

where $V_t(W_t, Z_t)$ equals the value function at time t and $u(W_T)$ the utility function of the wealth at time t, subject to a set of constraints. Equation 1 is a multi-period problem. It can be written to a single-period problem by using backward induction:

$$V_t(W_t, Z_t) = \max_{x_t} E_t \left[V_{t+1}(W_t(x_t^\top R_{e_{t+1}} + R_f), Z_{t+1}) \right] \tag{2}$$

In this paper, the assumption of a CRRA utility function is made:

$$U_T = \frac{W_T^{1-\gamma}}{1-\gamma}, \tag{3}$$

and therefore equation 1 can be simplified to (see Van Binsbergen & Brandt (2007) for a more detailed explanation):

$$
\begin{aligned}
V_t(W_t, Z_t) &= \max_{\{x_s\}_{s=t}^{T-1}} E_t \left[\left(\frac{W_t^{1-\gamma}}{1-\gamma} \right) \right] \\
&= \max_{\{x_s\}_{s=t}^{T-1}} E_t \left[\frac{W_t^{1-\gamma}}{1-\gamma} \prod_{s=t}^{T-1} (x_s' r_{s+1} + r_f)^{1-\gamma} \right] \\
&= \max_{x_t} E_t \left[\underbrace{\frac{W_t^{1-\gamma}(x_t^\top r_{t+1} + r_f)^{1-\gamma}}{1-\gamma}}_{u(W_{t+1})} \underbrace{\max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1} \left(\prod_{s=t+1}^{T-1} (x_s^\top r_{s+1} + r_f) \right)^{1-\gamma}}_{\psi(z_{t+1})} \right] \\
&= \max_{x_t} E_t \left[u(W_{t+1}) \psi(z_{t+1}) \right]. \qquad (4)
\end{aligned}
$$

Since the optimization is independent of wealth, W_T can be set to 1. Therefore equation 4 can be simplified to the Bellman Equation (Barron & Ishii (1989)):

$$
\frac{1}{1-\gamma} \psi(z_t) = \max_{x_t} E_t \left[u(x_t^\top r_{t+1} + r_f) \psi(z_{t+1}) \right]. \qquad (5)
$$

The term $\frac{1}{1-\gamma}$ is needed to get the first term in 4 to the form of $u(x_s^\top r_{t+1} + R_f)$. Therefore, to calculate the scaled value function, the right side of equation 5 needs to be multiplied with $(1-\gamma)$.

4.2 BGGS method

The BGGS method was introduced by Brandt et al. (2005). It is a different dynamic programming method, compared to discretizing the state space (DSS). An explanation of the DSS method can be found in A.1. The BGSS method can be explained in 5 steps based on Van Binsbergen & Brandt (2007):

- Step 1: Simulate N paths of simulated returns and simulated state variables of length T. The paths are based on an estimated VAR with the log excess returns and the state variable(s).
- Step 2: For every point in the N paths calculate the utility values using the CRRA utility.
- Step 3: Use the N simulated state variables to regress the N expected utilities from step 2 on these state variables. The fitted values of these regressions represent the expected utilities.
- Step 4: To calculate the utility in Step 2 we have to use a certain weight. Since the optimal weight is not known yet, we repeat steps 2 and 3 for a certain grid of weights. In this paper, the assumption is made that an investor is not allowed to short-sale or leverage, and therefore the grid of the weights is between 0 and 1 in steps of 0.01. When obtaining the expected utilities for each weight grid, we pick the optimal weight for each path.
- Step 5: Repeat steps 2 until 4 for $T-2, T-3, \ldots\ldots, 0$ to obtain an optimal weight at time $t = 0$ for every path. By taking the average of the optimal weight for all the paths, the

weight at time $t = 0$ is determined. This analysis is repeated for a certain number of iterations. The final weight at time $t = 0$ is the mean over all iterations.

The method can be solved by either iterating on the optimal portfolio weights or on the value function. The difference between these two approaches is explained in A.3.

Van Binsbergen & Brandt (2007) conclude that iterating on the value function will lead to more biased portfolio weights. In this paper, we will replicate their results and base our choice of iterating either on the value function or portfolio weights based on these replication results.

4.3 Portfolio Construction

To research the real-world performance of the BGSS method, different portfolios will be constructed to examine the out-of-sample portfolio performance. An investor can allocate his wealth between a risky asset, namely the S&P 500 and the risk-free rate, and has the option to rebalance his/her portfolio every quarter. The return of this portfolio can be computed as:

$$r_t^P = \mathbf{w}_t^T \mathbf{r}_t, \tag{6}$$

where $\mathbf{w}_t^T$ is the vector of weights at quarter t, $\mathbf{r}_t$ is the vector with returns at quarter t and N is the number of assets in the portfolio (in this paper $N = 2$).

For the BGSS method, we will examine the impact of using different state variables mentioned in Section 3 on portfolio performance. At every point in time, the BGSS method will be used for the remainder of the investment horizon. The simulations are conducted with an unrestricted VAR model where a rolling window is used for the estimations.

The general considered benchmark model is the $1/N$ portfolio. In this strategy, every asset is equally weighted, and therefore $w_t^{(i)}$ is equal to $1/N$ for every asset i at every point in time t (in this paper N=2).

To evaluate if taking future expectations into account is beneficial, another benchmark model is used where at each point in time we will only look one period into the future. This means the BGSS method is used without the regressions. In other words, at every point in time, the returns are simulated one period ahead, and the weight is based on these simulated returns. In this method, only the future expectations of one period ahead are taken into account therefore allowing for a comparison against the BGSS method where the future expectations of the whole investment period are taken into account.

This paper assumes a constant risk-free rate in line with the methodology presented in Van Binsbergen & Brandt (2007). At every point in time, the risk-free rate is updated to the most recent risk-free rate. However, for an investment horizon of $T = 20$ and $T = 40$, the risk-free rate is set constant on the mean of the last 10 years at every point in time. The rationale behind this approach is that for a long investment period, the risk-free rate is less likely to stay at the most recent risk-free rate in comparison to a short investment period. Therefore, for these two long investment periods, we set the risk-free rate at every point in time equal to the 10 year previous mean. For the benchmark models, the risk-free rate is always set to the most recent risk-free rate, also for investment horizons 20 and 40. This approach is appropriate because these models always only look one period ahead, making the most recent risk-free rate a suitable choice.

After this, more state variables are used simultaneously. The results are compared with the performance of the portfolios constructed on the single state variables and against the benchmark $1/N$ model. This analysis will help us determine if incorporating more state variables improves portfolio performance.

All the methods are used for all the investment periods mentioned in Section 3 and for four different levels of risk aversion ($\gamma = 5, 10, 15, 20$). The shorter the investment period the more often the analysis needs to be performed for more reliable results. For example, when $T = 2$, one model can provide much better results than another model for 2023Q3-2023Q4 whereas the reverse could be true for 2022Q3-2022Q4. For the investment periods $T = 2, 4, 8, 12, 20, 40$, the analysis is performed 20,10,5,3,2,1 times, respectively. This approach ensures all the investment periods fall into the investment period mentioned in Section 3 (2014Q1 until 2023Q3) and each iteration of the analysis is performed on a new investment period that does not contain any overlapping quarters with previous investment periods. For example for $T = 20$, the first investment period is 2019Q1-2023Q4 and the second is 2014Q1-2018Q4. For $T = 2$, the first investment period is 2023Q3-2024Q4, the second is 2023Q1-2023Q2, and so forth. This might not be necessary when Newey-West standard errors are taken into account, but since we conduct the analysis with a limited number of iterations anyway, we chose for non-overlapping investment periods. The mean and standard deviation for the performance measures across all iterations are taken to compare the models. The number of iterations is still too low to draw reliable conclusions. Therefore it is important to note that one should perform the analysis more times to make sure the conclusions are robust. The results in this paper give a good insight into how the models perform, but they are not fully reliable due to the limited number of iterations.

4.4 Performance measures

Evaluating the real performance of the methods is done by using the actual return and variance of the portfolio, the Sharpe ratio, and the Certainty Equivalent.

The Sharpe ratio is defined as follows:

$$S = \frac{R_p - R_f}{\sigma_p}, \tag{7}$$

where R_p is the return of the portfolio, R_f is the risk-free rate (which means the numerator equals the excess return), and σ_p equals the standard deviation of the excess returns. The Sharpe ratio is a measure to evaluate the performance of a portfolio by adjusting for the portfolio its risk.

The Certainty Equivalent is based on the power utility. The power utility function is given by:

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}, \tag{8}$$

where W is the wealth (which can be set to 1) and γ is the coefficient of risk aversion. To calculate the Certainty Equivalent, we need the expected utility:

$$E[U(W)] = E\left[\frac{(1+X)^{1-\gamma}}{1-\gamma}\right], \tag{9}$$

where X represents the return on the investment. The CEQ can therefore be calculated as:

$$CEQ^p = (E[U(W)] \cdot (1-\gamma))^{\frac{1}{1-\gamma}} . \quad (10)$$

The Certainty Equivalent represents the return an investor would consider equivalent to the investment strategy. Similar to the Sharpe ratio, the CEQ measures the portfolio's risk-return trade-off.

5 Results

5.1 Portfolio vs Value iteration

First, to confirm the results from Van Binsbergen & Brandt (2007) that iterating on the value function gives more biases than iterating on the portfolio iteration, we replicate an empirical example they performed using both options of the BGSS method. The results of the optimal weights at time $t = 0$ are shown in Section B. Since the method is based on simulations of returns and state variables we do not obtain the exact same weights as Van Binsbergen & Brandt (2007) but the conclusions remain consistent. Using the DSS method as a benchmark, the value function results in more biased weights compared to iterating on the optimal portfolio weights, especially for a high level of risk aversion or a high investment horizon. For example, for $T = 12$ with a level of risk aversion of $\gamma = 20$, iterating on the value function gives (when $N = 1000$) a weight of 0.8652 while iterating on the portfolio weights gives 0.2672. If we compare these with the weight obtained using the DSS method (0.25), the weight from iterating on the value function is significantly more biased than the weight from iterating on the portfolio weights. As the investment horizon increases both methods lose some precision, but value iteration loses significantly more. For example, when $T = 20$ with a level of risk aversion of $\gamma = 10$ ($N = 1000$), iterating on the value function yields a weight of 0.8788, while a weight of 0.6355 is obtained when iterating on portfolio weights. Compared to 0.65 (weight obtained with DSS), iterating on portfolio weights results again in less biased outcomes. So, the conclusion from Van Binsbergen & Brandt (2007) that iterating on portfolio weights is superior to iterating on the value function is confirmed. Therefore, for the rest of this paper, the BGSS method will be used with portfolio iteration rather than value iteration.

The results in Van Binsbergen & Brandt (2007) are based on a VAR model estimated by Barberis (2000). Now, the results are based on an investment scenario where an investor can allocate his wealth between the S&P 500 and the risk-free rate. Therefore, the VAR model is estimated on the log excess return of the value-weighted S&P 500 index. At every point in time, the VAR is re-estimated using a rolling window. To compare the difference, the VAR model for the same period as estimated by Barberis (2000) is estimated and shown in A.5. The estimated coefficients look similar, as well as the covariance matrix of the residuals. The main difference is that we use an unrestricted VAR(1) model, whereas in Van Binsbergen & Brandt (2007) they use a restricted VAR(1) model where the dependent variables only depend on the lagged state variable and not on the lagged return.

Since the computation time can become substantial when simulating 80 iterations of $N =$

1000 paths at every point in time, the decision must be made to either short the number of simulations or the number of paths (N). If we look at the results of the replication, the standard deviations are generally higher with 100 simulated paths compared to 1000 simulated paths. Therefore, maintaining a high number of simulated paths is preferable to ensure more stable results. That is why the simulated paths are only reduced to 500 and not to 100. Additionally, the number of simulations is reduced to 60. While this may increase the errors in the model, it will significantly reduce the computation time.

To assess the impact of these adjustments, the empirical example on the VAR estimated by Barberis (2000) is performed with these parameter settings. The results are shown in Table 14. These results indicate a slight increase in standard errors but not drastically. For example, for $T = 2$ with $\gamma = 10$ the standard error is 0.0630, in comparison to 0.0494 for $N = 1000$ and 80 simulations, and for $T = 20$ with $\gamma = 15$ the standard error is 0.0899 compared to 0.0887.

For an investment period of $T = 40$, we had to reduce the parameters even further. Since the method has to be repeated at every point in the investment period the computation time when $T = 40$ quarters becomes large. Ultimately, the method is performed with 100 paths for an investment period of 40 quarters. This will impact the standard errors significantly, but it is beyond the scope of this research to implement the model with higher parameter settings due to the high computation time.

5.2 Results individual state variables

Table 3 shows the performance of the single state variable BGSS models compared to the $1/N$ model. The table shows that for an investment period of 8 quarters or longer, the LDP model consistently outperforms the $1/N$ model. For every investment period T with a length equal to or longer than 8 quarters, the LDP model has at least two levels of risk aversion for which the Sharpe ratio is higher. For high investment periods, $T = 20$ and $T = 40$, the LDP model provides clear better results. For example, for $T = 40$ with a level of risk aversion of $\gamma = 20$, the Sharpe ratio is 0.91 compared to 0.75 for the $1/N$ model, and also the other levels of risk aversion have higher Sharpe ratios than the $1/N$ model. However, it is important to note that the analysis is only implemented once for the investment period of 40 quarters. To confirm the robustness of the results, the analysis has to be performed more often. The Certainty Equivalent for the LDP model is higher for all investment periods than for the $1/N$ model. This means that the return an investor would consider equivalent to the investment strategy is, in general, higher for the BGSS model with LDP as the state variable.

The BGSS models with LTR and DYS as the state variables perform worse than the model with LDP as the state variable. They are not able to outperform the $1/N$ portfolio, but the differences are not that high. For a short investment period of 2 horizons the BGSS method with DYS as the state variable performs very well for $\gamma = 5$ and $\gamma = 10$. However, this is partly due to a few exceptionally high Sharpe ratios that significantly increase the mean. For example, for $\gamma = 10$ one of the obtained Sharpe ratios is 31.82. Since the investment period is so short, this can happen when the excess returns are high with a low variance. This also becomes clear if we look at the standard deviations (23.54) of the obtained Sharpe ratios. Overall, while these BGSS models perform similarly to the $1/N$ model, the $1/N$ model generally performs better.

For example, for $T = 12$ and $\gamma = 10$, the 1/N provides a Sharpe Ratio of 0.90 while the DYS and LTR models provide, respectively, a Sharpe Ratio of 0.87 and 0.82. However, the differences are not high. Besides, an advantage of these models against the $1/N$ model is the variance (shown in Table 7), which is systematically lower for $\gamma = 10, 15, 20$. Additionally, the standard errors in the Sharpe Ratios are generally lower for these BGSS models than for the 1/N model for higher levels of risk aversion. This means that the variation in Sharpe Ratio across different investment periods (with the same investment horizon T) is lower for these BGSS models which gives an investor more certainty. Although the Sharpe ratios might be lower, the reduction in variance and standard deviation of Sharpe Ratios is substantial, making these models potentially more appealing to risk-averse investors.

Table 3: Sharpe ratios/Certainty Equivalents individual

T	γ	LDP	DYS	LTR	$1/N$	LDP (%)	DYS (%)	LTR (%)	$1/N$ (%)
$T = 2$	5	2.31 (3.58)	11.09 (19.95)	6.46 (11.95)	6.31 (14.45)	8.75 (12.24)	10.92 (3.59)	6.03 (10.92)	6.66 (10.34)
	10	1.30 (1.39)	8.95 (23.54)	5.90 (15.77)	6.31 (14.45)	6.88 (10.92)	9.74 (2.95)	5.72 (9.54)	5.77 (11.19)
	15	1.36 (1.60)	3.68 (2.73)	1.92 (1.88)	6.31 (14.45)	5.75 (10.73)	9.06 (2.68)	5.02 (9.33)	5.12 (11.89)
	20	1.73 (2.72)	3.38 (2.37)	1.85 (1.97)	6.31 (14.45)	4.96 (10.85)	8.43 (2.52)	4.46 (9.37)	4.59 (12.43)
$T = 4$	5	1.95 (1.88)	1.89 (2.07)	1.78 (1.87)	2.01 (2.22)	8.84 (8.09)	5.36 (9.76)	5.16 (9.67)	6.61 (8.22)
	10	1.66 (1.40)	1.34 (1.26)	1.23 (1.18)	2.01 (2.22)	6.17 (6.32)	4.13 (7.53)	4.00 (6.83)	4.84 (3.61)
	15	1.41 (1.03)	1.13 (1.00)	1.05 (0.94)	2.01 (2.22)	4.88 (5.98)	3.47 (7.31)	3.38 (6.07)	3.61 (9.82)
	20	1.28 (0.88)	1.03 (0.87)	0.97 (0.84)	2.01 (2.22)	4.03 (6.02)	2.85 (6.59)	2.89 (5.96)	2.53 (10.78)
$T = 8$	5	1.65 (1.65)	1.24 (1.03)	1.05 (1.84)	1.39 (1.70)	8.36 (0.34)	5.79 (3.03)	6.21 (2.69)	5.61 (3.47)
	10	1.51 (1.22)	1.05 (0.72)	1.03 (0.82)	1.39 (1.70)	5.90 (2.05)	3.98 (1.88)	4.18 (1.88)	3.83 (3.52)
	15	1.32 (0.77)	0.96 (0.45)	1.13 (0.34)	1.39 (1.70)	4.81 (1.55)	2.84 (1.41)	2.89 (1.17)	1.92 (4.34)
	20	1.20 (0.54)	0.91 (0.38)	1.11 (0.53)	1.39 (1.70)	4.18 (1.27)	2.16 (1.00)	2.48 (0.99)	0.00 (0.06)
$T = 12$	5	1.02(0.26)	0.86 (0.63)	0.78 (0.56)	0.90 (0.53)	8.42 (3.42)	4.57 (0.68)	4.02 (1.21)	5.41 (0.54)
	10	1.04 (0.17)	0.87 (0.46)	0.82 (0.44)	0.90 (0.53)	5.91 (2.12)	3.56 (0.32)	2.72 (1.44)	3.39 (1.68)
	15	1.03 (0.17)	0.81 (0.30)	0.78 (0.30)	0.90 (0.53)	4.70 (1.74)	3.04 (0.46)	1.98 (1.14)	1.26 (0.03)
	20	1.04(0.18)	0.79 (0.24)	0.77 (0.28)	0.90 (0.53)	4.06 (1.31)	2.76 (0.50)	1.71 (0.96)	−0.88 (4.91)
$T = 20$	5	0.94 (0.16)	0.63 (0.00)	0.59 (0.06)	0.79 (0.04)	7.95 (1.17)	4.36 (1.17)	3.86 (0.75)	5.55 (2.07)
	10	0.81 (0.06)	0.44 (0.23)	0.39 (0.20)	0.79 (0.04)	4.75 (2.15)	2.16 (1.74)	1.81 (2.10)	3.67 (0.44)
	15	0.71 (0.22)	0.36 (0.35)	0.32 (0.32)	0.79 (0.04)	2.95 (2.79)	1.15 (2.19)	0.98 (2.89)	1.60 (1.30)
	20	0.63 (0.33)	0.39 (0.31)	0.37 (0.28)	0.79 (0.04)	1.54 (3.16)	0.42 (2.54)	0.29 (3.47)	−0.57 (3.03)
$T = 40$	5	0.85 (x)	0.69 (x)	0.63 (x)	0.75 (x)	7.72 (x)	5.07 (x)	4.55 (x)	5.53 (x)
	10	0.87 (x)	0.69 (x)	0.63 (x)	0.75 (x)	5.17 (x)	3.26 (x)	2.96 (x)	3.67 (x)
	15	0.88 (x)	0.73 (x)	0.68 (x)	0.75 (x)	3.92 (x)	2.76 (x)	2.55 (x)	1.58 (x)
	20	0.91 (x)	0.70 (x)	0.72 (x)	0.75 (x)	3.31 (x)	2.32 (x)	2.35 (x)	−0.70 (x)

Average and standard deviation (between brackets) of annualized Sharpe ratios and annualized Certainty Equivalents using the BGSS method with three different state variables: log dividend price ratio (LDP), default yield spread (DYS), and the return on long-term government bonds (LTR). $1/N$ corresponds with the $1/N$ benchmark model. The left part of the table shows the Sharpe Ratios and the right part the Certainty Equivalents. The performance measures are reported for risk aversion γ equal to 5,10,15 and 20 and for investment horizon T (in quarters) equal to 2,4,8,12,20 and 40.

To understand why the BGSS model with the LDP as the state variable performs better than the other models, we look at the weights obtained for each model. Generally, the LDP model provides higher weights than the other BGSS models and the $1/N$ model. As an example, we look at an investment period of 12 quarters with a level of risk aversion of 5. Table 3 shows the average results and from the table, we can see that for $T = 12$ and $\gamma = 5$ the LDP model performs the best with a Sharpe ratio of 1.02, compared to 0.86, 0.78, and 0.90 for the DYS, LTR and 1/N model, respectively. Figure 1 shows the obtained weights for $T = 12$ and $\gamma = 5$ for all three investment periods. From the figure, it becomes clear that in the second two investment periods, the LDP model provides significantly higher weights, and for the first investment period significantly lower weights. The figure shows that the LTR and DYS models provide much more similar and constant weights and that the LDP model deviates more.

This figure shows three plots of the assigned weights to the risky asset by the different BGSS models. The state variables used are the log dividend price ratio (LDP), long-term return on government bonds (LTR), and the default yield spread (DYS). The investment horizon is 12 quarters and the level of risk aversion is 5. The left plot corresponds with the weights for the period 2021Q1-2023Q4, the middle plot with the weights for the period 2018Q1-2020Q4, and the right plot with the weights for the period 2015Q1-2017Q4. The dashed black line is a straight line at 0.50 and corresponds with the $1/N$ model.

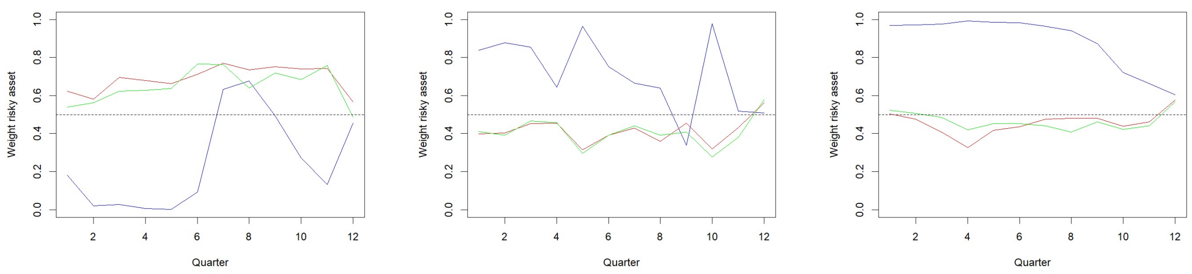

(a) 2021Q1-2023Q4 (b) 2018Q1-2020Q4 (c) 2015Q1-2017Q4

LDP	Blue line	**LTR**	Green line
DYS	Red line	**1/N**	Dashed black line

Figure 1: Weights assigned to the risky asset

To understand what the different weights result in, we look at the results for each investment period for all three models. Table 4 shows the results for each investment period. The mean and standard deviations of these periods give the results for $T = 12$ and $\gamma = 5$ in Tables 3 and 7. Period 1 corresponds with the first investment period (2021Q1-2023Q4) and plot 1a, period 2 with the second investment period (2018Q1-2020Q4) and plot 1b and period 3 (2015Q1-2017Q4) with the third investment period and plot 1c. Table 4 shows that in the second and third period (where the LDP provide higher weights) the mean returns are indeed the highest for this model. However, this also results in a higher variance. This makes sense as higher variance should be compensated with a higher mean in returns. If we look at the summary statistics of the S&P 500 (in Table 1) it can be observed that the mean of the log excess return is positive. Consequently, higher weights can lead to higher excess returns. As shown in Table 7, the mean returns are generally indeed the highest for the LDP. Although the variance is also higher, the mean returns are high enough to compensate for the higher variance and to therefore obtain a good Sharpe ratio. The reason why the model with LDP as the state variable still performs better for the first investment period than the other two models can be explained by the high variances obtained with the BGSS models with DYS and LTR as state variables, which are 34.17 and 33.69. Besides, the mean returns of the S&P 500 in the period 2021Q1-2023Q4 are 0.027 compared to 0.040 for 2018Q1-2020Q4 and 0.036 for 2015Q1-2018Q4. This shows that the lower weights do not result in much lower mean returns (which can be seen in Table 4), due to the weaker performance of the S&P 500 during that period. Moreover, in the years before 2021, the S&P 500 performed well and therefore the higher weights result in better performance. The mean returns for the DYS and LTR models are much more constant, but because the variance is much more volatile the Sharpe ratios vary significantly over the periods. Overall, the LDP model compensates the best for higher variance with higher mean return than the other two models, therefore obtaining better Sharpe ratios.

Table 4: Performance measures $T = 12$

	1	2	3
Sharpe	0.74	0.96	1.37
Mean	5.65	17.51	9.58
V	7.29	59.97	10.76
CEQ	4.89	11.82	8.46

(a) LDP

	1	2	3
Sharpe	0.47	0.53	1.58
Mean	7.58	6.72	5.57
V	34.17	22.02	2.76
CEQ	3.97	4.43	5.30

(b) DYS

	1	2	3
Sharpe	0.35	0.56	1.43
Mean	6.13	6.71	5.42
V	33.69	19.74	3.17
CEQ	2.52	4.68	5.11

(c) LTR

This table shows the results for three investment periods with a horizon of 12 quarters for three BGSS models with different state variables: log dividend price ratio (LDP), default yield spread (DYS), and the long-term return on government bonds (LTR). Period 1 corresponds with 2021Q1-2023Q4, period 2 with 2018Q1-2020Q4, and period 3 with 2015Q1-2017Q4. The table reports the annualized Sharpe Ratio (Sharpe), annualized mean returns of the portfolio (Mean), the quarterly variance of the portfolio (V), and the annualized Certainty Equivalent (CEQ).

We evaluate the simulated returns based on the VAR model to understand why the weights are significantly higher for some periods when using LDP as the state variable compared to the LTR and DYS. Therefore, we examine the difference between simulated returns in the first three quarters of the third investment period in plot 1c. In all these quarters, the LDP model generates higher weights for the risky asset. At these points in time the simulated paths are, respectively, of length 12,11,10, since the remainder of the investment horizon is, respectively, $T = 12, 11, 10$. For each state variable, the mean is taken of each simulated return for all the simulated paths. Then the average over the number of iterations (60) is taken. Table 5 shows the results of this. From the table, it becomes clear that the log dividend price ratio simulates higher returns. This partly explains the higher weights assigned by the BGSS model with LDP as the state variable. If the simulated return for the next period is higher, it is more likely that the fitted value from the regression is higher, and therefore the most optimal weight is higher. The DYS and LTR provide similar simulated returns and this explains why the weights in figure 1 move together. The results in Table 10 also support this. All results in this table are based on the benchmark models, where the weights are determined based on a return that is only simulated one period into the future. The table indicates that the LDP has the highest mean returns for almost all investment horizons and levels of risk aversion. This is again the fact of the higher weights assigned to the risky asset. If this model assigns higher weights it means the simulated return for one period into the future is more often positive. If the simulations increase to longer paths (equal to an investment horizon) as done in the BGSS method, the number of positive simulated returns can increase and therefore even higher weights are assigned to the risky asset.

Table 5: Mean of simulated log excess returns first period

	$T = 12$	$T = 11$	$T = 10$
LDP	0.020 (0.001)	0.020 (0.001)	0.020 (0.001)
DYS	0.017 (0.001)	0.017 (0.001)	0.017 (0.001)
LTR	0.017 (0.001)	0.017 (0.001)	0.017 (0.001)

Mean and standard deviation (between brackets) of simulated log excess returns based on a path of length T (remanding investment horizon) on a VAR(1) model between the log excess return and a certain state variable: log dividend price ratio (LDP), default yield spread (DYS) or the return on longer-term government bonds (LTR). This table shows the average of 60 simulations. The original investment horizon is $T = 20$ from 2015Q1-2017Q4.

Another example is given in Table 6. However, this is for the time period where the LDP assigns significantly lower weights to the risky asset. Here, the simulated returns are on average significantly lower for the LDP model therefore explaining the lower weights assigned to the risky asset. This is probably due to COVID-19 which happened in the first quarter of 2020. The first investment period in Figure 1 starts at 2021Q1, which means the estimation period includes 2020 and therefore contains COVID-19. From the figure and Table 6, the conclusion can be drawn that the LDP model adjust significantly more after COVID-19 since the simulated returns drop substantially more in comparison to the other two models. Plot 1b also supports this, since the LDP weights drop significantly from the 10th to 11th quarter which is right after COVID-19 started. Plot 1c shows that after 7 quarters the first higher weight above 0.50 is back, but also the 11th quarter has a low weight again. It will be interesting to research in the future if the model keeps assigning low weights to the risky asset or if it turns around and starts assigning higher weights to the risky asset again.

Table 6: Mean of simulated log excess returns third period

	$T = 12$	$T = 11$	$T = 10$
LDP	0.005 (0.001)	0.004 (0.001)	0.003 (0.001)
DYS	0.017 (0.001)	0.017 (0.001)	0.017 (0.001)
LTR	0.017 (0.001)	0.017 (0.001)	0.017 (0.001)

Mean of and standard deviation (between brackets) of simulated log excess returns based on a path of length T (remanding investment horizon) on a VAR(1) model between the log excess return and a certain state variable: log dividend price ratio (LDP), default yield spread (DYS) or the return on long-term government bonds (LTR). This table shows the average of 60 simulations. The original investment horizon is $T = 20$ from 2021Q1-2023Q4.

The results in Table 7 show the property of the power utility: the lower the risk aversion, the higher the mean returns and the higher the variance. For all investment periods, the mean returns and variance decline when moving to a higher level of risk aversion. Therefore, investors can make optimal decisions based on their own risk aversion by adjusting the level of risk aversion. Table 3 shows that for most investment periods the Sharpe Ratios of the BGSS models are the highest for a lower level of risk aversion ($\gamma = 5$). This indicates that these models perform, on average, better for this level of risk aversion. This is interesting, as the results might improve even more if the level of risk aversion would be even lower. The levels of risk aversions are inspired from Van Binsbergen & Brandt (2007), but the results indicate that a higher level of risk aversion seems to work better for investment periods $T = 2$ until $T = 20$. As expected, the Certainty Equivalents decrease when the level of risk aversion is increased.

Table 7: Mean returns/variance individual

T	γ	LDP	DYS	LTR	1/N	LDP	DYS	LTR	1/N
$T = 2$	5	9.44 (10.30)	9.16 (12.77)	7.63 (10.28)	7.49 (9.60)	21.65 (39.29)	16.29 (22.32)	20.85 (43.79)	19.04 (44.27)
	10	8.14 (8.82)	8.93 (11.23)	7.23 (9.24)	7.49 (9.60)	14.79 (25.33)	17.78 (19.15)	15.52 (27.17)	19.04 (44.27)
	15	7.65 (8.59)	8.31 (10.11)	7.06 (9.09)	7.49 (9.60)	13.42 (22.38)	17.29 (16.92)	14.47 (22.77)	19.04 (44.27)
	20	7.46 (8.67)	8.14 (9.89)	7.01 (9.16)	7.49 (9.60)	12.78 (20.25)	17.59 (16.19)	14.03 (20,89)	19.04 (44.27)
$T = 4$	5	10.47 (9.71)	6.89 (9.31)	6.41 (9.49)	7.38 (8.02)	20.71 (33.24)	16.18 (18.67)	15.76 (17.22)	15.18 (22.40)
	10	7.81 (7.17)	5.61 (6.31)	5.39 (6.55)	7.38 (8.02)	10.64 (14.27)	8.92 (8.50)	8.91 (7.85)	15.18 (22.40)
	15	6.64 (6.04)	5.18 (5.34)	5.04 (5.55)	7.38 (8.02)	7.59 (7.93)	7.43 (6.41)	7.41 (5.99)	15.18 (22.40)
	20	6.03 (5.46)	4.94 (5.10)	4.88 (5.18)	7.38 (8.02)	6.53 (5.98)	7.05 (5.86)	6.88 (5.29)	15.18 (22.40)
$T = 8$	5	10.50 (3.28)	6.16 (3.17)	6.39 (3.92)	7.21 (3.94)	24.40 (28.39)	17.41 (6.75)	20.83 (18.52)	16.09 (13.84)
	10	8.05 (2.60)	5.03 (1.79)	4.02 (1.77)	7.21 (3.94)	13.98 (19.61)	7.41 (3.04)	6.47 (4.37)	16.09 (13.84)
	15	6.66 (1.88)	4.71 (1.30)	4.09 (1.16)	7.21 (3.94)	8.38 (10.67)	4.81 (1.88)	4.54 (2.06)	16.09 (13.84)
	20	5.86 (1.52)	4.44 (1.08)	4.32 (1.13)	7.21 (3.94)	5.73 (5.76)	3.51 (1.67)	3.56 (1.70)	16.09 (13.84)
$T = 12$	5	10.92 (4.93)	6.62 (1.00)	6.09 (0.14)	7.24 (1.66)	26.01 (24.06)	19.65 (15.84)	18.20 (5.80)	17.69 (13.95)
	10	8.02 (3.36)	4.89 (1.00)	4.37 (1.02)	7.24 (1.66)	11.60 (9.40)	6.87 (5.11)	6.51 (2.63)	17.69 (13.95)
	15	6.45 (2.55)	4.30 (1.22)	4.16 (0.94)	7.24 (1.66)	6.57 (4.35)	4.77 (3.17)	4.96 (1.18)	17.69 (13.95)
	20	5.67 (1.97)	4.53 (1.78)	3.97 (0.99)	7.24 (1.66)	4.72 (2.61)	4.06 (2.52)	3.86 (1.13)	17.69 (13.95)
$T = 20$	5	11.07 (3.52)	6.20 (3.25)	5.74 (2.37)	7.19 (3.50)	29.58 (22.85)	16.66 (14.25)	16.82 (14.23)	15.19 (13.08)
	10	8.74 (3.02)	3.64 (3.03)	3.31 (2.68)	7.19 (3.50)	20.93 (16.00)	6.64 (3.24)	6.61 (3.14)	15.19 (13.08)
	15	6.84 (3.42)	2.75 (3.01)	2.57 (2.79)	7.19 (3.50)	13.89 (10.07)	4.60 (0.97)	4.44 (0.78)	15.19 (13.08)
	20	5.52 (3.49)	2.37 (2.93)	2.18 (2.77)	7.19 (3.50)	9.79 (4.89)	3.88 (0.27)	3.68 (0.00)	15.19 (13.08)
$T = 40$	5	10.74 (x)	6.78 (x)	6.64 (x)	7.17 (x)	28.94 (x)	16.90 (x)	16.41 (x)	15.15 (x)
	10	9.38 (x)	4.47 (x)	4.72 (x)	7.17 (x)	20.64 (x)	6.58 (x)	6.42 (x)	15.15(x)
	15	8.25 (x)	3.90 (x)	3.28 (x)	7.17 (x)	15.00 (x)	3.77 (x)	3.17 (x)	15.15 (x)
	20	7.55 (x)	3.73 (x)	3.03 (x)	7.17 (x)	11.55 (x)	3.08 (x)	2.44 (x)	15.15 (x)

Average and standard deviation (between brackets) of annualized mean returns and quarterly variance using the BGSS method with three different state variables: log dividend price ratio (LDP), default yield spread (DYS), and the return on long-term government bonds (LTR). $1/N$ corresponds with the $1/N$ benchmark model. The left part of the table shows the returns and the right part the variance. The performance measures are reported for risk aversion γ equal to 5,10,15 and 20 and for investment horizon T (in quarters) equal to 2,4,8,12,20 and 40.

The tables in A.6 show the performance of the benchmark models obtained from the state variables. So, each weight is obtained by only simulating one period into the future. The results show that these models look similar to the 1/N model. The weights of these benchmark models for most points in time all fluctuate between 0.40 and 0.60, which explains the similar results. From these results and the results from the individual BGSS models, the conclusion can be drawn that taking future expected utility into account is not always profitable since the BGSS models do not always beat their benchmark model. Especially for the BGSS models with LTR and DYS as the state variable, their benchmark model often performs better. The BGSS model with LDP as the state variable generally performs better than the benchmark model, especially for a longer investment period or a higher level of risk aversion.

The advantage of different returns and variances for different levels of risk aversions disappears for these benchmark models, since for all investment periods the difference in performance between the different levels of risk aversion is non-significant. This makes sense since we only simulate a log excess return for one period and whenever this is positive the optimal weight will be 1 and whether this is negative it will be 0 since in both cases this will maximize the expected utility. So, the level of risk aversion does not make a difference here.

5.3 Results Multiple State Variables

Table 8 shows the results of the BGSS method with multiple state variables. Generally, the combinations LDP/LTR and LDP/DYS outperform the $1/N$ model based on their Sharpe ratios. These models provide higher Sharpe ratios, particularly for higher investment periods (8 quarters or more). However, the LTR/DYS consistently underperforms compared to the $1/N$ model. As shown in Table 3, the LDP model performs significantly better than the LTR and DYS models individually. Therefore it is not surprising that the combinations including LDP perform better than the LTR/DYS combination. Among all models the LDP/DYS performs the best, often outperforming the single LDP model based on the Sharpe ratio. For example, for $T = 20$ the LDP/LTR model provides a Sharpe ratio of 0.95 for $\gamma = 5$ and 0.86 for $\gamma = 10$. The LDP model provides a Sharpe ratio of, respectively, 0.94 and 0.81. However, the differences in the Sharpe ratios are small. Also the differences with the certainty Equivalents obtained with the single LDP model are small. Compared to the $1/N$ model, these combination models, as well as the individual LDP model, often obtain higher Certainty Equivalents which favors these models over the $1/N$ model. The mean returns and variances are shown in Table 11. The mean returns are indeed the highest for the two combinations with the LDP and comparable to those of the individual LDP models as shown in Table 7.

Table 8: Sharpe Ratios and Certainty Equivalents combination models

T	γ	LTR/DYS	LDP/DYS	LDP/LTR	All	$1/N$	LTR/DYS (%)	LDP/DYS (%)	LDP/LTR(%)	All(%)	$1/N$ (%)
$T = 2$	5	5.35 (9.17)	3.81 (5.64)	3.04 (3.38)	6.79 (17.11)	6.31 (14.45)	6.60 (10.69)	9.79 (10.90)	8.40 (9.98)	8.33 (2.95)	6.66 (10.34)
	10	4.81 (9.76)	4.05 (6.76)	2.58 (3.44)	3.23 (10.23)	6.31 (14.45)	5.69 (9.37)	7.85 (9.72)	6.65 (8.48)	6.66 (4.67)	5.77 (11.19)
	15	2.02 (2.14)	3.10 (3.01)	4.53 (12.32)	2.64 (2.61)	6.31 (14.45)	5.03 (9.22)	7.61 (9.33)	5.84 (8.26)	7.26 (6.25)	5.12 (11.89)
	20	1.81 (1.82)	3.74 (6.72)	2.34 (2.86)	9.03 (33.97)	6.31 (14.45)	4.53 (9.16)	6.48 (9.18)	5.17 (8.16)	6.86 (6.91)	4.59 (12.43)
$T = 4$	5	1.80 (1.85)	1.58 (2.32)	1.93 (1.91)	2.54 (2.47)	2.01 (2.22)	5.56 (10.09)	8.62 (8.59)	8.63 (8.14)	8.73 (1.97)	6.13 (2.22)
	10	1.26 (1.19)	1.68 (1.90)	1.69 (1.44)	2.33 (2.03)	2.01 (2.22)	4.18 (7.00)	5.23 (8.09)	6.13 (6.47)	6.71 (2.83)	4.84 (8.87)
	15	1.08 (0.97)	1.40 (1.44)	1.42 (1.09)	1.93 (1.71)	2.01 (2.22)	3.47 (6.23)	5.61 (7.99)	4.70 (5.94)	6.51 (2.83)	3.61 (9.81)
	20	0.99 (0.87)	1.15 (1.28)	1.30 (0.93)	1.88 (1.62)	2.01 (2.22)	2.94 (6.16)	4.82 (7.64)	3.94 (6.11)	5.49 (2.77)	2.53 (10.78)
$T = 8$	5	1.22 (1.22)	1.58 (1.66)	1.66 (1.74)	1.16 (1.92)	1.39 (1.70)	4.83 (4.96)	7.01 (4.57)	8.25 (3.02)	7.10 (4.84)	5.62 (3.47)
	10	1.04 (0.69)	1.49 (1.36)	1.54 (1.24)	1.29 (1.50)	1.39 (1.70)	3.79 (2.48)	5.27 (3.36)	6.09 (1.76)	5.43 (3.65)	3.82 (3.52)
	15	0.96 (0.50)	1.35 (1.09)	1.32 (0.80)	1.14 (0.49)	1.39 (1.70)	3.33 (1.58)	4.60 (2.91)	4.91 (1.34)	3.80 (2.52)	1.91 (4.33)
	20	0.91 (0.41)	1.42 (0.80)	1.18 (0.59)	1.02 (0.50)	1.39 (1.70)	3.04(1.24)	4.00(2.49)	4.06(1.12)	2.17 (1.24)	0.00 (5.63)
$T = 12$	5	0.80 (0.31)	1.13 (0.28)	1.02 (0.29)	1.12 (0.38)	0.90 (0.53)	4.30 (1.48)	7.47 (1.90)	5.97 (1.71)	7.43 (1.75)	5.41 (0.54)
	10	0.82 (0.16)	1.12 (0.22)	1.03 (0.23)	1.15 (0.23)	0.90 (0.53)	3.49 (1.28)	5.20 (1.62)	5.09 (1.62)	5.57 (0.97)	3.34 (1.61)
	15	0.79 (0.08)	1.08 (0.20)	1.06 (0.20)	1.11 (0.19)	0.90 (0.53)	2.99 (1.15)	4.65 (1.44)	4.66 (1.53)	4.34 (0.38)	1.26 (3.27)
	20	0.76 (0.19)	1.08 (0.19)	1.08 (0.18)	1.10 (0.17)	0.90 (0.53)	2.64 (1.08)	4.30 (1.36)	4.33 (1.46)	3.85 (0.21)	-0.01 (4.91)
$T = 20$	5	0.62 (0.05)	0.95 (0.19)	0.94 (0.09)	0.93 (0.21)	0.76 (0.01)	3.98 (1.11)	7.87 (0.35)	5.30 (1.17)	7.89 (0.09)	4.67 (0.18)
	10	0.40 (0.19)	0.86 (0.04)	0.80 (0.04)	0.85 (0.07)	0.76 (0.01)	2.64 (1.34)	4.78 (0.72)	4.68 (1.11)	4.61 (1.20)	3.16 (0.62)
	15	0.33 (0.31)	0.77 (0.12)	0.71 (0.16)	0.75 (0.06)	0.76 (0.01)	2.17 (1.63)	2.67 (1.16)	3.04 (2.12)	2.51 (0.98)	1.49 (1.42)
	20	0.30 (0.37)	0.69 (0.24)	0.64 (0.25)	0.67 (0.16)	0.79 (0.01)	1.88 (1.88)	1.33 (1.69)	1.68 (2.75)	1.09 (1.53)	0.00 (0.02)
$T = 40$	5	0.68 (x)	0.86 (x)	0.88 (x)	0.88 (x)	0.75 (x)	3.77 (x)	7.48 (x)	7.84 (x)	7.61 (x)	5.55 (x)
	10	0.66 (x)	0.92 (x)	0.88 (x)	0.88 (x)	0.75 (x)	3.51 (x)	5.03 (x)	5.37 (x)	4.92 (x)	3.67 (x)
	15	0.66 (x)	0.91 (x)	0.89 (x)	0.92 (x)	0.75 (x)	3.19 (x)	3.41 (x)	4.00 (x)	3.66 (x)	1.58 (x)
	20	0.69 (x)	0.97 (x)	0.93 (x)	0.92 (x)	0.75 (x)	2.82 (x)	2.93 (x)	3.43 (x)	2.64 (x)	-0.01 (x)

Average and standard deviation (between brackets) of annualized Sharpe ratios and annualized Certainty Equivalents using the benchmark BGSS method where the weight is obtained by only looking one period into the future with three different state variables: log dividend price ratio (LDP), default yield spread (DYS) and the return on long-term government bonds (LTR). $1/N$ corresponds with the $1/N$ benchmark model. The left part of the table shows the Sharpe ratios and on the right the Certainty Equivalents. The Sharpe ratios are reported for risk aversion γ equal to 5,10,15 and for investment horizon T (in quarters) equal to 2,4,8,12,20 and 40.

The 'All' columns corresponds with the performance of the BGSS model using all three state variables simultaneously. The results show that the model provides similar results to the two combinations including LDP. For lower investment horizons, this combination performs better, for example, obtaining a Sharpe ratio of 9.03 for $T = 2$ and $\gamma = 20$ and of 2.54 for $T = 4$ and $\gamma = 5$, which are higher than for other combinations. However, the standard deviation of the Sharpe ratios therefore increases significantly.

Overall, all combinations with the LDP perform better than all individual models. Although the differences with the individual LDP model are small, the combinations with LDP generally provide slightly higher Sharpe ratios. In most cases, the combinations with the LDP provide a slightly higher Sharpe ratio. Nonetheless, there are also situations where the opposite is true as well. The combination LTR/DYS performs similarly to the two individual models and significantly worse than the combinations including LDP. Therefore, using combinations of state variables can be beneficial but the state variables must perform well individually to achieve good results.

Figure 2 displays the assigned weights to the risky asset by the different combination models over an investment period of 20 quarters and a level of risk aversion of $\gamma = 10$. The plots show both investment periods of 20 quarters. The plots show that, except for the quarters after COVID-19, all three combination models including the LDP assign significantly higher weights to the risky asset than the LTR/DYS model. Overall, similar to the individual models, the models with the LDP generally allocate higher weights to the risky asset, but they demonstrate more downfall post COVID-19. Again, the higher weights lead to higher mean returns and higher variance for the LDP models, as shown in Table 11.

Figure 2: Weights assigned to the risky asset by combination models

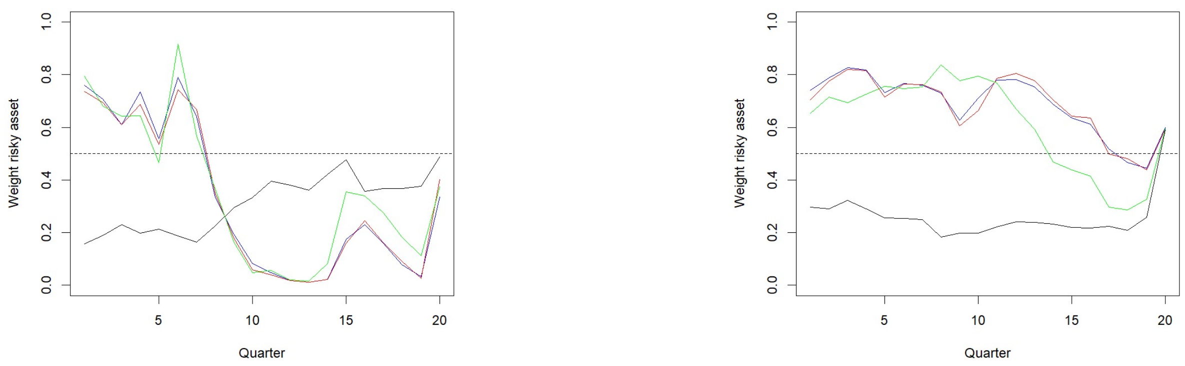

(a) 2019Q1-2023Q4

(b) 2014Q1-2018Q4

All	Blue line	**LDP/DYS**	Red line
LDP/LTR	Green line	**LTR/DYS**	Black line
1/N	Dashed black line		

This figure shows two plots of the assigned weights to the risky asset by the different BGSS models. The different state variables used are the log dividend price ratio (LDP), the return on long-term government bonds (LTR), and the default yield spread (DYS). 'All' corresponds with the combination where all three state variables are used simultaneously. The investment period equals 20 quarters and the level of risk aversion equals 10. The left plot corresponds with the weights for the period 2019Q1-2023Q4 and the right plot for the period 2014Q1-2018Q4. The dashed black line is a straight line at 0.50 and corresponds with the $1/N$ model.

6 Conclusion

This paper evaluates the out-of-sample performance of the BGSS method, proposed by Brandt et al. (2005). This evaluation is conducted by implementing the method in a real-world investment scenario where an investor allocates his wealth between a risky asset (the S&P 500) and the risk-free rate, for a multi-period investing problem: investing from time t until T with the option to rebalance the portfolio at each point in between t and $T-1$.

The BGSS method is assessed using three different state variables: the log dividend price ratio (LDP), the return on long-term government bonds (LTR), and the default yield spread (DYS). Additionally, to investigate if using more state variables simultaneously, a potential advantage over other dynamic programming methods such as discretizing the state space, improves portfolio performance we apply the BGSS method with multiple state variables simultaneously.

The replication of Van Binsbergen & Brandt (2007) shows that iterating on portfolio weights is superior to iterating on the value function, as it results in significantly less biased portfolio weights. Therefore, the BGSS method is used with portfolio iteration. Individually, the BGSS model with the log dividend price ratio (LDP) as the state variable consistently outperforms the $1/N$ model for an investment horizon of 8 quarters or longer. For all these investment horizons, the model has at least two levels of risk aversion which provides a higher Sharpe ratio. Besides, it consistently provides higher Certainty Equivalents, indicating that the return an investor considers equal is higher than for the $1/N$ model. The BGSS models with the default yield spread (DYS) and the return on long-term government bonds (LTR) as the state variable perform significantly worse and did not consistently outperform the 1/N model. The LDP model consistently provides higher Sharpe ratios than the other models. This can be attributed to the generally higher assigned weights to the risky asset in the years before 2021 where the S&P 500 performed well due to its higher mean returns compared to the returns in the period 2021-2023. The higher weights are partially explained due to the on average higher simulated returns by the LDP model. Although the $1/N$ model provides better Sharpe ratios, the LTR and DYS models show potential, especially for risk-averse investors. This is due to the variance of the portfolio, which is consistently lower than for the 1/N model for high levels of risk aversion. Moreover, the variation in Sharpe Ratios is lower. This means it provides more stable Sharpe ratios for different investment periods of the same length. This can be appealing for risk-averse investors since it gives more certainty. All models generally perform better for a lower level of risk aversion, as the best Sharpe ratios are obtained with a level of risk aversion of $\gamma = 5$. This makes one wonder if the results could improve further if the level of risk aversion is decreased further.

For the combinations, the results increase slightly for the combinations including the LDP. However, these improvements are not significant. The combination with all three state variables did not improve the results compared to the combinations with two state variables including the LDP. The combination with LTR and DYS as state variables performs significantly worse, which is not surprising when looking at the individual results. This shows that combinations of state variables can improve performance slightly, but careful selection of state variables is crucial for obtaining optimal portfolio performance.

This research is limited in several aspects. Therefore, there is a lot of room for further research. First, in the investment scenario, it was only possible to invest in one risky asset, namely the S&P 500. Considering more possible risky assets could give a deeper understanding of how the BGSS method performs in different investment scenarios. Moreover, it can be interesting to research how the method would perform on financial instruments that have non-linear pay-offs, such as options. Besides, the research was limited to three different state variables, although more variables could be good candidates.

We used the method similarly as in Van Binsbergen & Brandt (2007). One of the assumptions made in the method is a constant risk-free rate. In this paper, the risk-free rate is updated at every point in time to the latest risk-free rate for investment horizons $T = 2, 4, 8, 12$ or set at the mean of the last 10 years for $T = 20, 40$, but then the assumption of a constant risk-free rate is used. Especially for long investment horizons this does not hold. This could impact the effectiveness of the model. For example, let the initial risk-free rate be set to 1.015 for the whole investment period. If in reality, the average risk-free rate in this period is 1.010 the weights assigned to the risky asset might be too low, and therefore higher weights could have been more optimal. Therefore evaluating how the model performs if we predict the risk-free rate instead of assuming it to stay constant for the whole investment period could provide a deeper understanding of the method.

Moreover, the new literature mentioned in 2 on this topic has not been implemented. Since the BGSS was first introduced, there has been done much more research which extended the method. Comparing the BGSS method to other solutions of dynamic programming problems could give more valuable insights.

Lastly, this paper only looks at a quarterly level of rebalancing and does not compare different levels of rebalancing. It can be interesting to research if increasing or decreasing the rebalancing frequency can improve the results and if the BGSS method is more suitable for a different rebalancing frequency.

References

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

Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance*, *55*(1), 225–264.

Barron, E. & Ishii, H. (1989). The bellman equation for minimizing the maximum cost. *Nonlinear Analysis, Theory, Methods and Applications*, *13*(9), 1067–1090.

Berry, W. D. & Feldman, S. (1985). *Multiple regression in practice.* Sage Publications.

Brandt, M. W. (1999). Estimating portfolio and consumption choice: A conditional euler equations approach. *The Journal of Finance*, *54*(5), 1609–1645.

Brandt, M. W., Goyal, A., Santa-Clara, P. & Stroud, J. R. (2005). A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *The Review of Financial Studies*, *18*(3), 831–873.

Cong, F. & Oosterlee, C. W. (2017). Accurate and robust numerical methods for the dynamic portfolio management problem. *Computational Economics*, *49*, 433–458.

DeMiguel, V., Garlappi, L. & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, *22*(5), 1915–1953.

Garlappi, L. & Skoulakis, G. (2009). Numerical solutions to dynamic portfolio problems: The case for value function iteration using taylor approximation. *Computational Economics*, *33*, 193–207.

Jain, S. & Oosterlee, C. W. (2015). The stochastic grid bundling method: Efficient pricing of bermudan options and their greeks. *Applied Mathematics and Computation*, *269*, 412–431.

Van Binsbergen, J. H. & Brandt, M. W. (2007). Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Computational Economics*, *29*, 355–367.

Welch, I. & Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, *21*(4), 1455–1508.

A Appendix

A.1 Discrete state space method

This method serves as a benchmark for the BGSS. In Van Binsbergen & Brandt (2007) it is used as a benchmark to compare the weights from the two BGGS methods. It can be explained in 7 steps:

1. Define a set of grid points of the state variables.
2. We start with the optimization problem at $T-1$. For each grid point, simulate N returns and determine the weight for which the expected utility is the highest.
3. Calculate the scaled value function based on the optimal weight determined in the second step for each grid point.
4. For $T-2$, simulate N returns and N state variables for each grid point.
5. By using interpolation, determine the scaled value function for each of the simulated state variables.
6. Determine the optimal weight for each of the simulated grid points and calculate the scaled value functions.
7. Repeat step 4 to 6 for $T-3$ until we reach t.

The weight at time $t = 0$ will than be decided on the actual value of the state variables. The optimal weight determined for the grid point which corresponds with the actual value will be the value of the final weight weight at time $t = 0$.

A.2 Histogram of Log Excess returns

Figure 3: Histogram of the quartely returns

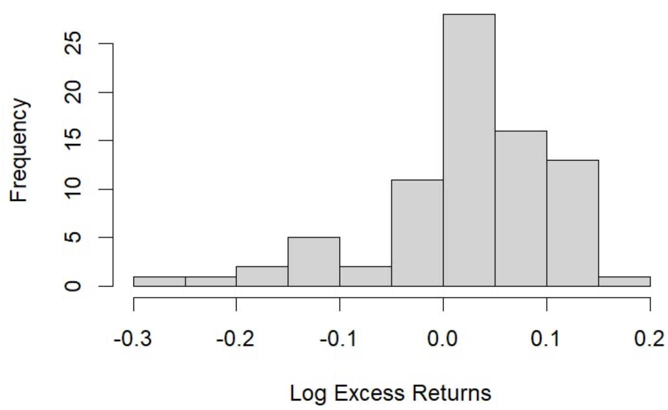

This graph shows the histogram of the quarterly log excess returns of the S&P 500. The data sample starts from 1994Q1 and ends at 2023Q4.

A.3 Difference between iterating on Portfolio weights or on the Value function

To show the difference between iterating on portfolio weights or on the value function, we take a look at the optimization problem at time $T-1$:

$$\max_{x_{T-1}} E_{T-1}\left(u(x_{T-1}^{\top}\mathrm{R^e}_T + Rf)\right), \tag{11}$$

For which the optimal x_{t-1}^{*} is stored for each simulation path and obtain the value function:

$$\frac{1}{1-\gamma}\psi_{T-1}(Z_{T-1}) = E_{T-1}\left(u({x^*}_{T-1}^{\top}\mathrm{R^e}_T + Rf)\right). \tag{12}$$

When using portfolio iteration the optimal x_{T-1}^{*} is stored and solve for $T-2$:

$$\max_{x_{T-2}} E_{T-1}\left[\frac{[(x_{T-2}^{\top}R_{T-1}^{e} + Rf)(x_{T-1}^{\top}R_{T}^{e} + Rf)]}{1-\gamma}\right].$$

However, when iterating on the value function, we do not store the optimal weights but the value function obtained with the optimal weights and in solve in period $T-2$:

$$\begin{aligned}&\max_{x_{T-2}} E_{T-2}\left[u(x_{T-2}^{\top}R_{T-1}^{e} + Rf)\psi_{T-1}(Z_{T-1})\right]\\ &= \max_{x_{T-2}} E_{T-2}\left[\left(x_{T-2}^{\top}R_{T-1}^{e} + Rf\right)^{1-\gamma} E_{T-1}\left(\frac{\left(x_{T-1}^{\top}R_{T}^{e} + Rf\right)^{\gamma}}{1-\gamma}\right)\right].\end{aligned}$$

Van Binsbergen & Brandt (2007) conclude that iterating on the value function will lead to more biased portfolio weights. The reason they give is that when portfolio weights are bound by, for example, short sale constraints the error is bounded as well.

Imagine there is an approximation error in the value function at time $T-1$. The value function is a conditional expectation itself, so an error here means we did not accurately estimate the expected utilities based on state variables. This is because a polynomial approximation is used, which is not fully accurate. In other words, this polynomial approximation is just an approximation of the real relationship between the state variables and the conditional moments. These errors in the value function can therefore lead to significant errors in the chosen weights.

However, with short sale constraints, the range of possible portfolio weights is limited, so the error is also limited. By iterating on portfolio weights, we keep the errors from growing too much over time. In contrast, if we iterate on the value function, we don't have this limit, and errors can become more significant.

A.4 Estimated VAR model Van Binsbergen & Brandt (2007)

$$\begin{pmatrix} r_{e,t+1} \\ d_{t+1}-p_{t+1}\end{pmatrix} = \begin{pmatrix} 0.227\,(0.95) \\ -0.155\,(-0.79)\end{pmatrix} + \begin{pmatrix} 0.060\,(0.87) \\ 0.958\,(17.02)\end{pmatrix}(d_t - p_t) + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1}\end{pmatrix} \tag{13}$$

$$\begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1}\end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0\end{pmatrix}, \begin{pmatrix} 0.0060 & -0.0051 \\ -0.0051 & 0.0049\end{pmatrix}\right), \tag{14}$$

where r^e_{t+1} is the log excess return on the value-weighted CRSP index at time $t+1$, $d_{t+1}-p_{t+1}$ is the log dividend price ratio at time $t+1$ computed from the sum of the past 12 monthly dividends and the current level of the index. They assume the log dividend price ratio is log (0.03), which equals the unconditional mean.

A.5 Example Estimated VAR Model

$$\begin{pmatrix} r_{e,t+1} \\ d_{t+1}-p_{t+1} \end{pmatrix} = \begin{pmatrix} 0.28955\,(0.924) \\ -0.23066\,(-0.755) \end{pmatrix} + \begin{pmatrix} -0.03735\,(-0.222) & 0.07841\,(0.861) \\ -0.06040\,(-0.368) & 0.93574\,(10.534) \end{pmatrix} \begin{pmatrix} r_{e,t} \\ d_t - p_t \end{pmatrix} + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix} \tag{15}$$

$$\begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.005937 & -0.005559 \\ -0.005559 & 0.005644 \end{pmatrix} \right), \tag{16}$$

where r^e_t is the log excess return on the value-weighted S&P 500 index at time t, $d_t - p_t$ is the log dividend price ratio at time t computed from the sum of the past 12 monthly dividends and the current level of the index. The VAR is estimated on quarterly data with an estimation period of 1986Q1-1995Q4.

Correlation Matrix of Residuals

The correlation matrix of residuals is given by:

$$\begin{pmatrix} \text{logexcessquar} & \text{quarterlydiv} \\ 1.0000 & -0.9603 \\ -0.9603 & 1.0000 \end{pmatrix}$$

A.6 Results individual benchmark

Table 9: Sharpe ratios/Certainty Equivalents individual Benchmarks

T	γ	LDP	DYS	LTR	$1/N$	LDP (%)	DYS (%)	LTR (%)	$1/N$ (%)
$T = 2$	5	7.19 (11.80)	6.04 (13.74)	6.78 (15.74)	6.31 (14.45)	7.40 (10.35)	7.20 (11.99)	7.11 (11.72)	6.66 (10.34)
	10	6.17 (10.26)	6.87 (16.60)	5.82 (12.16)	6.31 (14.45)	6.41 (11.38)	6.23 (12.98)	6.12 (12.69)	5.77 (11.19)
	15	6.59 (10.67)	6.25 (14.33)	6.92 (16.61)	6.31 (14.45)	5.63 (12.28)	5.45 (13.80)	5.39 (13.48)	5.12 (11.89)
	20	6.94 (11.17)	6.20 (14.20)	5.65 (11.40)	6.31 (14.45)	5.03 (12.87)	4.81 (14.40)	4.78 (14.06)	4.59 (12.43)
$T = 4$	5	2.02 (2.06)	2.00 (2.21)	2.01 (2.10)	2.01 (2.22)	7.03 (8.39)	6.65 (9.57)	6.61 (9.06)	6.61 (8.22)
	10	2.02 (2.06)	2.00 (2.20)	2.01 (2.11)	2.01 (2.22)	5.39 (9.15)	5.01 (10.49)	4.97 (9.90)	4.84 (3.61)
	15	2.01 (2.05)	2.00 (2.21)	2.00 (2.11)	2.01 (2.22)	3.94 (10.22)	3.46 (11.69)	3.39 (11.09)	3.61 (9.82)
	20	2.02 (2.06)	2.00 (2.22)	2.00 (2.10)	2.01 (2.22)	2.65 (11.30)	2.15 (12.80)	2.12 (12.11)	2.53 (10.78)
$T = 8$	5	1.42 (1.55)	1.41 (1.55)	1.40 (1.53)	1.39 (1.70)	6.24 (3.63)	5.87 (3.69)	5.91 (3.71)	5.61 (3.47)
	10	1.42 (1.55)	1.40 (1.55)	1.39 (1.52)	1.39 (1.70)	3.88 (3.87)	3.51 (4.01)	3.53 (4.05)	3.83 (3.52)
	15	1.42 (1.55)	1.40 (1.53)	1.40 (1.54)	1.39 (1.70)	1.39 (5.03)	1.08 (5.17)	1.09 (5.14)	1.92 (4.34)
	20	1.42 (1.55)	1.40 (1.54)	1.40 (1.53)	1.39 (1.70)	-1.10 (6.56)	-1.14 (6.71)	-1.13 (6.71)	0.00 (0.92)
$T = 12$	5	0.91 (0.43)	0.89 (0.44)	0.88 (0.44)	0.90 (0.53)	5.85 (0.67)	5.57 (0.61)	5.54 (0.61)	5.41 (0.54)
	10	0.91 (0.43)	0.89 (0.44)	0.89 (0.44)	0.90 (0.53)	3.10 (2.06)	2.91 (1.98)	2.95 (1.96)	3.39 (1.68)
	15	0.91 (0.43)	0.89 (0.44)	0.89 (0.44)	0.90 (0.53)	0.24 (3.82)	0.13 (3.66)	0.15 (3.67)	1.26 (0.03)
	20	0.90 (0.42)	0.88 (0.43)	0.89 (0.44)	0.90 (0.53)	-2.54 (5.55)	-2.62 (5.26)	-2.51 (5.23)	-0.88 (4.91)
$T = 20$	5	0.79 (0.02)	0.79 (0.04)	0.79 (0.04)	0.79 (0.04)	5.89 (1.43)	5.77 (1.31)	5.78 (1.31)	5.55 (2.07)
	10	0.79 (0.02)	0.78 (0.03)	0.79 (0.03)	0.79 (0.04)	3.22 (0.23)	3.23 (0.25)	3.25 (0.25)	3.67 (0.44)
	15	0.79 (0.03)	0.79 (0.04)	0.79 (0.04)	0.79 (0.04)	0.34 (2.01)	0.54 (1.93)	0.53 (1.86)	1.60 (1.20)
	20	0.79 (0.03)	0.79 (0.04)	0.79 (0.04)	0.79 (0.04)	-2.58 (3.54)	-2.22 (3.42)	-2.31 (3.41)	-0.57 (3.03)
$T = 40$	5	0.74 (x)	0.74 (x)	0.74 (x)	0.75 (x)	5.83 (x)	5.81 (x)	5.83 (x)	5.53 (x)
	10	0.74 (x)	0.75 (x)	0.75 (x)	0.75 (x)	3.18 (x)	3.20 (x)	3.21 (x)	3.67 (x)
	15	0.74 (x)	0.75 (x)	0.74 (x)	0.75 (x)	0.33 (x)	0.51 (x)	0.52 (x)	1.58 (x)
	20	0.74 (x)	0.74 (x)	0.75 (x)	0.75 (x)	-2.62 (x)	-2.27 (x)	-2.30 (x)	-0.70 (x)

Average and standard deviation (between brackets) of annualized Sharpe ratios and annualized Certainty Equivalents using the benchmark BGSS method where the weight is obtained by only looking one period into the future with three different state variables: log dividend price ratio (LDP), default yield spread (DYS) and the return on long-term government bonds (LTR). $1/N$ corresponds with the $1/N$ benchmark model. The left part of the table shows the Sharpe ratios and the right the Certainty Equivalents. The Sharpe ratios are reported for risk aversion γ equal to 5,10,15 and 20 and for investment horizon T (in quarters) equal to 2,4,8,12,20 and 40.

Table 10: Mean returns/Variance individual Benchmarks

T	γ	LDP	DYS	LTR	$1/N$	LDP	DYS	LTR	$1/N$
$T = 2$	5	8.23 (12.25)	8.35 (11.06)	8.27 (10.80)	7.49 (9.60)	25.69 (64.72)	23.80 (52.73)	23.69 (52.71)	19.04 (44.27)
	10	8.29 (12.23)	8.35 (11.06)	8.28 (10.77)	7.49 (9.60)	25.57 (64.19)	23.82 (52.72)	23.86 (53.04)	19.04 (44.27)
	15	8.25 (12.29)	8.36 (11.07)	8.28 (10.80)	7.49 (9.60)	25.50 (63.94)	23.98 (53.46)	23.66 (52.46)	19.04 (44.27)
	20	8.28 (12.25)	8.36 (11.07)	8.29 (10.80)	7.49 (9.60)	25.64 (64.71)	24.06 (53.58)	23.74 (52.53)	19.04 (44.27)
$T = 4$	5	9.76 (7.11)	8.24 (9.24)	8.20 (8.74)	7.38 (8.02)	19.63 (30.94)	19.35 (27.26)	19.27 (25.79)	15.18 (22.40)
	10	9.75 (7.12)	8.25 (9.25)	8.19 (8.72)	7.38 (8.02)	19.63 (31.00)	19.31 (27.24)	19.24 (25.77)	15.18 (22.40)
	15	9.83 (7.11)	8.20 (9.25)	8.16 (8.71)	7.38 (8.02)	19.60 (31.76)	19.28 (27.04)	19.38 (26.08)	15.18 (22.40)
	20	9.76 (7.11)	8.21 (9.24)	8.18 (8.71)	7.38 (8.02)	19.67 (31.23)	19.33 (27.24)	19.31 (25.95)	15.18 (22.40)
$T = 8$	5	8.39 (4.15)	7.93 (1.35)	7.98 (4.08)	7.21 (3.94)	21.50 (18.49)	20.82 (17.80)	20.61 (15.80)	16.09 (13.84)
	10	8.38 (4.13)	7.93 (1.35)	7.96 (4.10)	7.21 (3.94)	21.25 (18.17)	20.64 (17.73)	20.65 (15.80)	16.09 (13.84)
	15	8.41 (4.17)	7.95 (1.38)	7.96 (4.07)	7.21 (3.94)	21.36 (18.39)	20.64 (17.68)	20.56 (15.74)	16.09 (13.84)
	20	8.38 (4.14)	7.94 (1.37)	7.96 (4.06)	7.21 (3.94)	21.29 (18.18)	20.67 (17.64)	20.55 (15.71)	16.09 (13.84)
$T = 12$	5	8.34 (1.62)	7.93 (1.35)	7.92 (1.36)	7.24 (1.66)	23.90 (18.5)	22.68 (18.46)	22.75 (18.39)	17.69 (13.95)
	10	8.34 (1.63)	7.93 (1.35)	7.97 (1.39)	7.24 (1.66)	23.87 (18.80)	22.68 (18.63)	22.70 (18.23)	17.69 (13.95)
	15	8.37 (1.68)	7.95 (1.38)	7.94 (1.35)	7.24 (1.66)	23.95 (18.87)	22.76 (18.50)	22.63 (18.26)	17.69 (13.95)
	20	8.36 (1.66)	7.94 (1.37)	7.96 (1.37)	7.24 (1.66)	23.94 (18.90)	22.74 (18.35)	22.60 (18.16)	17.69 (13.95)
$T = 20$	5	7.93 (1.35)	7.92 (2.66)	7.94 (2.65)	7.19 (3.50)	20.77 (18.12)	19.51 (12.00)	19.52 (12.01)	15.19 (13.08)
	10	7.92 (1.35)	7.60 (2.67)	7.93 (2.68)	7.19 (3.50)	20.77 (18.28)	19.61 (12.10)	19.64 (12.13)	15.19 (13.08)
	15	7.95 (1.38)	7.93 (2.64)	7.93 (2.66)	7.19 (3.50)	20.77(18.26)	19.57 (12.11)	19.55 (12.03)	15.19 (13.08)
	20	7.94 (1.37)	7.93 (2.64)	7.90 (2.64)	7.19 (3.50)	20.68 (18.05)	19.44 (12.00)	19.56 (12.00)	15.19 (13.08)
$T = 40$	5	8.11 (x)	7.93 (x)	7.90 (x)	7.17 (x)	20.70 (x)	19.44 (x)	19.31 (x)	15.15 (x)
	10	8.10 (x)	7.91 (x)	7.99 (x)	7.17 (x)	20.82 (x)	19.30 (x)	19.49 (x)	15.15 (x)
	15	8.15 (x)	7.96 (x)	7.90 (x)	7.17 (x)	20.82 (x)	19.37 (x)	19.31 (x)	15.15 (x)
	20	8.08 (x)	7.88 (x)	7.99 (x)	7.17 (x)	20.50 (x)	19.24 (x)	19.58 (x)	15.15 (x)

Average and standard deviation (between brackets) of annualized mean returns and quarterly variance using the benchmark BGSS method where each weight is obtained by only looking one period into the future with three different state variables: log dividend price ratio (LDP), default yield spread (DYS) and the return on long-term government bonds (LTR). $1/N$ corresponds with the $1/N$ benchmark model. The left part of the table shows the returns and the right part the variance. The performance measures are reported for risk aversion γ equal to 5,10,15 and 20 and for investment horizon T (in quarters) equal to 2,4,8,12,20 and 40.

A.7 Mean returns/variance combinations BGSS models

Table 11: Mean returns/Variance Combinations

T	γ	DYS/LTR	LDP/DYS	LDP/LTR	All	$1/N$	DYS/LTR (%)	LDP/DYS (%)	LDP/LTR (%)	All (%)	$1/N$ (%)
$T=2$	5	8.90 (3.79)	7.97 (3.38)	11.84 (12.23)	10.99 (9.37)	7.49 (9.60)	25.75 (49.43)	30.03 (80.83)	22.55 (47.52)	25.01 (80.76)	19.04 (44.27)
	10	8.79 (2.75)	8.48 (2.93)	11.93 (9.75)	8.41 (8.05)	7.49 (9.60)	18.96 (29.88)	16.74 (50.87)	13.19 (28.53)	15.03 (60.46)	19.04 (44.27)
	15	8.11 (2.23)	8.48 (2.70)	11.68 (9.22)	7.01 (9.27)	7.49 (9.60)	16.31 (24.32)	13.83 (43.50)	12.11 (28.62)	14.78 (49.80)	19.04 (44.27)
	20	8.78 (2.33)	8.44 (2.60)	11.62 (8.80)	7.00 (9.14)	7.49 (9.60)	14.83 (22.51)	12.72 (38.16)	10.77 (27.17)	15.02 (46.72)	19.04 (44.27)
$T=4$	5	7.73 (9.72)	10.69 (7.34)	11.18 (10.50)	11.33 (8.78)	7.38 (8.02)	14.18 (19.44)	20.18 (18.80)	20.25 (24.89)	18.27 (40.89)	15.18 (22.40)
	10	5.43 (4.66)	7.58 (6.12)	8.65 (6.38)	7.60 (5.75)	7.38 (8.02)	9.86 (7.25)	10.67 (6.96)	10.32 (5.61)	11.00 (18.00)	15.18 (22.40)
	15	5.25 (3.69)	6.60 (5.13)	7.87 (4.77)	6.69 (4.88)	7.38 (8.02)	8.90 (5.14)	7.82 (4.28)	7.32 (3.31)	7.65 (8.72)	15.18 (22.40)
	20	4.96 (2.92)	5.83 (4.66)	7.24 (4.01)	6.02 (4.28)	7.38 (8.02)	7.43 (4.67)	6.96 (3.36)	6.31 (2.24)	7.73 (6.79)	15.18 (22.40)
$T=8$	5	6.57 (4.09)	10.01 (3.02)	10.68 (3.77)	9.11 (3.33)	7.21 (3.94)	17.72 (16.54)	23.10 (24.00)	23.21 (19.90)	22.24 (24.64)	16.09 (13.84)
	10	5.10 (2.00)	7.94 (2.17)	8.03 (3.38)	7.62 (2.22)	7.21 (3.94)	7.03 (6.09)	11.20 (11.35)	12.55 (6.53)	10.70 (12.42)	16.09 (13.84)
	15	4.59 (1.34)	6.71 (1.58)	6.70 (2.03)	6.58 (1.60)	7.21 (3.94)	4.85 (3.53)	7.15 (6.32)	8.15 (4.74)	6.83 (6.26)	16.09 (13.84)
	20	4.36 (1.18)	5.97 (1.31)	5.42 (2.00)	5.50 (1.41)	7.21 (3.94)	4.11 (2.68)	5.45 (3.94)	5.81 (3.24)	5.25 (3.14)	16.09 (13.84)
$T=12$	5	6.40 (1.32)	9.85 (3.24)	10.75 (5.08)	9.79 (4.37)	7.24 (1.66)	20.02 (13.60)	23.02 (22.60)	24.15 (22.02)	22.96 (28.22)	17.69 (13.95)
	10	4.87 (1.50)	7.70 (2.29)	8.07 (3.79)	7.71 (3.18)	7.24 (1.66)	7.21 (4.32)	10.84 (8.92)	11.58 (9.85)	10.69 (11.55)	17.69 (13.95)
	15	4.25 (1.40)	6.33 (1.69)	6.50 (2.48)	6.23 (2.17)	7.24 (1.66)	4.72 (2.48)	6.57 (4.57)	6.39 (4.61)	6.36 (5.99)	17.69 (13.95)
	20	3.99 (0.97)	5.64 (1.28)	5.75 (2.10)	5.61 (1.75)	7.24 (1.66)	3.93 (1.88)	4.59 (2.54)	4.56 (2.79)	4.59 (3.64)	17.69 (13.95)
$T=20$	5	6.10 (1.98)	10.86 (1.72)	10.86 (2.42)	10.65 (2.08)	7.19 (3.50)	17.24 (10.27)	27.10 (12.83)	27.40 (13.98)	27.36 (18.32)	15.19 (13.08)
	10	3.46 (1.95)	8.51 (2.02)	8.37 (3.01)	8.50 (2.92)	7.19 (3.50)	6.87 (2.49)	17.43 (7.33)	19.31 (10.19)	18.42 (11.86)	15.19 (13.08)
	15	2.46 (1.99)	6.62 (2.16)	6.71 (3.16)	6.73 (3.24)	7.19 (3.50)	4.58 (0.71)	11.58 (4.15)	13.80 (6.50)	12.76 (7.73)	15.19 (13.08)
	20	2.32 (1.99)	5.51 (2.24)	5.51 (3.00)	5.60 (3.32)	7.19 (3.50)	3.83 (0.13)	8.74 (2.67)	9.47 (3.93)	9.50 (4.76)	15.19 (13.08)
$T=40$	5	6.76 (x)	10.48 (x)	10.43 (x)	10.30 (x)	7.17 (x)	16.24 (x)	25.66 (x)	24.23 (x)	24.80 (x)	15.15 (x)
	10	4.57 (x)	9.03 (x)	8.99 (x)	8.79 (x)	7.17 (x)	6.54 (x)	17.84 (x)	18.31 (x)	17.17 (x)	15.15 (x)
	15	3.81 (x)	7.77 (x)	7.86 (x)	7.93 (x)	7.17 (x)	4.10 (x)	12.44 (x)	13.11 (x)	12.51 (x)	15.15 (x)
	20	3.66 (x)	7.15 (x)	7.19 (x)	7.15 (x)	7.17 (x)	3.39 (x)	9.02 (x)	9.91 (x)	9.88 (x)	15.15 (x)

Average and standard deviation (between brackets) of annualized mean returns and quarterly variance (standard deviation between brackets) using the BGSS method with combinations of the state variables: log dividend price ratio (LDP), default yield spread (DYS), and the return on long-term government bonds (LTR). The left part of the table shows the returns and on the right part the variance. The performance measures are reported for risk aversion γ equal to 5,10,15 and 20 and for investment horizon T (in quarters) equal to 2,4,8,12,20 and 40.

B Results Replication

Table 12: Weight in Stocks using Portfolio Weight Iteration

Horizon (T)	γ	$N = 100$	$N = 1000$	DSS
$T = 2$	5	0.6371 (0.2058)	0.6468 (0.0938)	0.65
	10	0.3389 (0.1335)	0.3257 (0.0494)	0.30
	15	0.2419 (0.1004)	0.2224 (0.0295)	0.24
	20	0.1968 (0.0825)	0.1636 (0.0211)	0.14
$T = 4$	5	0.6111 (0.1962)	0.7002 (0.0826)	0.71
	10	0.3624 (0.1610)	0.3625 (0.0427)	0.34
	15	0.2889 (0.1168)	0.2381 (0.0259)	0.29
	20	0.2321 (0.0848)	0.1782 (0.0217)	0.18
$T = 8$	5	0.7431 (0.1959)	0.7645 (0.0865)	0.96
	10	0.4708 (0.1802)	0.4385 (0.0522)	0.48
	15	0.3539 (0.1465)	0.2853 (0.0431)	0.28
	20	0.2794 (0.1275)	0.2210 (0.0291)	0.22
$T = 12$	5	0.7411 (0.2291)	0.8521 (0.0694)	0.87
	10	0.5325 (0.1848)	0.5079 (0.0694)	0.43
	15	0.4095 (0.1615)	0.3497 (0.0599)	0.35
	20	0.3672 (0.1269)	0.2672 (0.0535)	0.25
$T = 20$	5	0.7971 (0.1806)	0.9255 (0.0564)	1.00
	10	0.6478 (0.1691)	0.6355 (0.0762)	0.65
	15	0.5259 (0.1635)	0.4624 (0.0887)	0.46
	20	0.4216 (0.1444)	0.3384 (0.0665)	0.38
$T = 40$	5	0.8723 (0.1342)	0.9676 (0.0327)	1.00
	10	0.6868 (0.2121)	0.7630 (0.0981)	0.85
	15	0.5991 (0.1805)	0.5345 (0.1248)	0.75
	20	0.4726 (0.1771)	0.4435 (0.1330)	0.57

Portfolio weight in the risky asset (S&P 500) at time $t = 0$ obtained by the BGSS method with the log dividend price ratio as state variable using portfolio function iteration, for risk aversion γ equal to 5, 10, 15, and 20, for investment horizon T (in quarters) equal to 2, 4, 8, 12, 20, and 40 and for a number of simulated paths N equal to 100 and 1000. The table reports the averages and standard deviations (between brackets) of 80 simulations. DSS corresponds with the obtained weights by the DSS method and serves as the benchmark.

Table 13: Weight in Stocks using Value function Iteration

Horizon (T)	γ	$N = 100$	$N = 1,000$	DSS
$T = 2$	5	0.6354 (0.1858)	0.6915 (0.0798)	0.65
	10	0.3354 (0.1327)	0.3202 (0.0425)	0.30
	15	0.2531 (0.1037)	0.2185 (0.0264)	0.24
	20	0.1922 (0.0868)	0.1656 (0.0217)	0.14
$T = 4$	5	0.6592 (0.1972)	0.6915 (0.0798)	0.71
	10	0.3913 (0.1257)	0.3566 (0.0407)	0.34
	15	0.3097 (0.1369)	0.2430 (0.0294)	0.29
	20	0.2998 (0.1680)	0.1864 (0.0292)	0.18
$T = 8$	5	0.7458 (0.1895)	0.7897 (0.0793)	0.96
	10	0.5420 (0.1435)	0.4435 (0.0533)	0.48
	15	0.5386 (0.2290)	0.3143 (0.0563)	0.28
	20	0.7148 (0.2781)	0.3678 (0.2259)	0.22
$T = 12$	5	0.7772 (0.1891)	0.8575 (0.0724)	0.87
	10	0.6463 (0.1494)	0.5344 (0.0603)	0.43
	15	0.7958 (0.2206)	0.5265 (0.1821)	0.35
	20	0.8836 (0.1937)	0.8652 (0.2304)	0.25
$T = 20$	5	0.8672 (0.1281)	0.9422 (0.0537)	1.00
	10	0.9282 (0.0858)	0.8788 (0.0790)	0.65
	15	0.8431 (0.2391)	0.9794 (0.0682)	0.46
	20	0.8067 (0.2374)	0.9387 (0.1375)	0.38
$T = 40$	5	0.9478 (0.0744)	0.9885 (0.0151)	1.00
	10	0.6099 (0.3145)	0.7683 (0.2021)	0.95
	15	0.8423 (0.2275)	0.9917 (0.0284)	0.75
	20	0.8279 (0.1833)	0.9824 (0.0481)	0.57

Portfolio weight in the risky asset (S&P 500) at time $t = 0$ obtained by the BGSS method with the log dividend price ratio as state variable using value function iteration, for risk aversion γ equal to 5, 10, 15, and 20, for investment horizon T (in quarters) equal to 2, 4, 8, 12, 20, and 40 and for a number of simulated paths N equal to 100 and 1000. The table reports the averages and standard deviations (between brackets) of 80 simulations. DSS corresponds with the obtained weights by the DSS method and serves as the benchmark.

Table 14: Weight in stocks using portfolio iteration with 500 paths and 60 iterations

Horizon(T)	γ	$N = 500$
$T = 2$	5	0.6391 (0.0946)
	10	0.3186 (0.0630)
	15	0.2305 (0.0406)
	20	0.1711 (0.0286)
$T = 4$	5	0.7178 (0.1088)
	10	0.3797 (0.0820)
	15	0.2603 (0.0485)
	20	0.1946 (0.0435)
$T = 8$	5	0.7639 (0.1211)
	10	0.4461 (0.0867)
	15	0.0.3014 (0.0660)
	20	0.2280 (0.0675)
$T = 12$	5	0.8354 (0.1016)
	10	0.5442 (0.0925)
	15	0.3589 (0.0671)
	20	0.2734 (0.0700)
$T = 20$	5	0.8933 (0.0841)
	10	0.6548 (0.1137)
	15	0.4751 (0.0899)
	20	0.3632 (0.0971)
$T = 40$	5	0.9506 (0.0513)
	10	0.7515 (0.1246)
	15	0.5368 (0.1527)
	20	0.4422 (0.1308)

Portfolio weight using portfolio function iteration in the risky asset (S&P 500) at time $t = 0$ obtained by the BGSS method with the log dividend price ratio as the state variable, for risk aversion γ equal to 5, 10, 15, and 20, for investment horizon T (in quarters) equal to 2, 4, 8, 12, 20, and 40 and for a number of simulated paths N equal to 500. The table reports the averages and standard deviations (between brackets) of 60 simulations.

Table 15: Certainty Equivalent using Portfolio iteration

Horizon (T)	γ	$N = 100$ (%)	$N = 1,000$ (%)
$T = 2$	5	10.04 (2.38)	8.95 (0.74)
	10	8.14 (1.23)	7.57 (0.39)
	15	7.43 (1.00)	7.10 (0.26)
	20	7.14 (0.63)	6.88 (0.20)
$T = 4$	5	11.59 (1.17)	10.68 (0.59)
	10	9.59 (0.91)	8.59 (0.37)
	15	8.53 (0.84)	7.78 (0.36)
	20	7.95 (0.77)	7.43 (0.18)
$T = 8$	5	11.69 (1.70)	11.25 (0.40)
	10	10.05 (0.91)	7.07 (0.30)
	15	9.05 (0.65)	8.16 (0.18)
	20	8.48 (0.64)	7.68 (0.18)
$T = 12$	5	11.90 (0.77)	11.38 (0.32)
	10	10.29 (0.75)	9.34 (0.24)
	15	9.35 (0.53)	8.42 (0.19)
	20	8.66 (0.56)	7.85 (0.16)
$T = 20$	5	11.87 (0.52)	11.40 (0.18)
	10	10.42 (0.52)	9.76 (0.20)
	15	9.53 (0.52)	8.75 (0.19)
	20	8.97 (0.51)	8.16 (0.18)
$T = 40$	5	11.44 (0.28)	11.04 (0.08)
	10	10.28 (0.29)	9.89 (0.13)
	15	9.65 (0.30)	9.12 (0.14)
	20	9.14 (0.27)	8.61 (0.19)

Annualized Certainty Equivalent obtained by the BGSS method with log dividend price ratio as the state variable using portfolio iteration, for risk aversion γ equal to 5, 10, 15, and 20, for investment horizon T (in quarters) equal to 2, 4, 8, 12, 20, and 40 and for a number of simulations of paths N of 100 and 1,000. The table reports the averages and standard deviations (between brackets) of 80 simulations.

Table 16: Certainty Equivalent using Value function iteration

Horizon (T)	γ	$N = 100(\%)$	$N = 1,000(\%)$
$T = 2$	5	9.95 (1.81)	9.19 (0.80)
	10	8.33 (1.29)	7.57 (0.36)
	15	7.36 (0.89)	7.09 (0.21)
	20	7.12 (0.76)	6.83 (0.17)
$T = 4$	5	11.12 (1.66)	10.63 (0.65)
	10	9.31 (1.45)	8.52 (0.33)
	15	8.58 (0.85)	7.77 (0.25)
	20	7.89 (0.75)	7.42 (0.21)
$T = 8$	5	12.14 (1.06)	11.19 (0.38)
	10	10.21 (1.00)	9.11 (0.25)
	15	9.25 (0.78)	8.22 (0.24)
	20	8.91 (0.88)	7.96 (0.45)
$T = 12$	5	12.02 (0.77)	11.40 (0.29)
	10	10.48 (0.77)	9.45 (0.22)
	15	9.91 (0.88)	8.68 (0.37)
	20	9.65 (0.89)	8.80 (0.54)
$T = 20$	5	11.96 (0.48)	11.44 (0.20)
	10	10.80 (0.51)	10.11 (0.21)
	15	10.17 (0.57)	9.80 (0.30)
	20	9.72 (0.55)	9.88 (0.37)
$T = 40$	5	11.52 (0.28	11.08 (0.08)
	10	9.68 (0.42)	10.09 (0.14)
	15	9.21 (0.37)	9.71 (0.17)
	20	9.15 (0.32)	9.75 (0.18)

Annualized Certainty Equivalent obtained by the BGSS method with log dividend price ratio as state variable using value iteration, for risk aversion γ equal to 5, 10, 15, and 20, for investment horizon T (in quarters) equal to 2, 4, 8, 12, 20, and 40 and for a number of simulations of paths N of 100 and 1,000. The table reports the averages and standard deviations (between brackets) of 80 simulations.

C Programming Code

In the provided zip-file all code and datasets can be found to replicate all results. The README file provided in the zip-file explains how the codes should be run and which results each file produces. As mentioned in the README file, further instructions (on how to switch between state variables or combinations) are mentioned in the comments of the R files. All computations in this paper are performed in RStudio version 2023.09.1. All needed R packages are at the top of each file.