

Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function?

Jules H. van Binsbergen · Michael W. Brandt

Received: 17 October 2006 / Accepted: 30 October 2006
© Springer Science+Business Media B.V. 2006

Abstract Most dynamic programming methods deployed in the portfolio choice literature involve recursions on an approximated value function. The simulation-based method proposed recently by Brandt, Goyal, Santa-Clara, and Stroud (*Review of Financial Studies*, 18, 831–873, 2005), relies instead on recursive uses of approximated optimal portfolio weights. We examine the relative numerical performance of these two approaches. We show that when portfolio weights are constrained by short sale restrictions for example, iterating on optimized portfolio weights leads to superior results. Value function iterations result in a lower variance but disproportionately higher bias of the solution, especially when risk aversion is high and the investment horizon is long.

Keywords Dynamic portfolio choice · Simulation method

1 Introduction

Recent progress in dynamic programming techniques has substantially enlarged the set of computationally feasible portfolio choice problems, including problems featuring model and parameter uncertainty, learning, background risks, and frictions. By far the most popular approach to solving dynamic portfolio choice problems starts by discretizing the state space, as done by Balduzzi

J. H. van Binsbergen (✉) · M. W. Brandt
Fuqua School of Business, Duke University, Durham, NC 27708, USA
e-mail: jules.vanbinsbergen@duke.edu

M. W. Brandt
NBER, Cambridge, MA, 10016 USA
e-mail: mbrandt@duke.edu

and Lynch (1999), Brandt (1999), Barberis (2000), and Dammon, Spatt, and Zhang (2001), among many others. Once the state space is discretized, the value function can be approximated by quadrature integration (Balduzzi and Lynch), simulations (Barberis), binomial discretization (Damon, Spatt and Zhang) or nonparametric regressions (Brandt). The dynamic program is then solved by standard backward recursion on the approximated value function. Specifically, each period, going backward from the second-to-last period to the first, the solution is found by maximizing the one-step ahead expectation of the approximated value function derived in the previous recursive step. In contrast to this common approach of recursing on the approximated value function, the simulation-based method proposed recently by Brandt, Goyal, Santa-Clara, and Stroud (BGSS) (2005), relies on recursive uses of approximated optimal portfolio weights to maximized expected terminal utility in every recursion. We examine the relative numerical performance of these two approaches. We show that when portfolio weights are constrained by short sale restrictions, for example, iterating on optimized portfolio weights leads to superior results. Value function iterations result in a lower variance but disproportionately higher bias of the solution, especially when risk aversion is high and the investment horizon is long.

The likely explanation of our results is the following. By iterating on approximated portfolio weights, which are bounded by the constraints, the approximation error that accumulates from one period to the next is also bounded. More importantly, as Cochrane (1989) points out, for standard utility functions first-order deviations in optimal policies lead to only second-order deviations in maximized utility (the value function). In other words, the expected utility function is very flat surrounding the optimal solution. As a result, even relatively large approximation errors in the optimized portfolio weights are associated with only small distortions of the optimization criterion. In contrast, in the value function iteration approach, approximation errors in the value function are not only first order but also accumulate further through the backward recursions.

The paper proceeds as follows. In Sect. 2, we describe the BGSS method of iterating on portfolio weights and explain further how it differs from iterating on the value function. Section 3 describes the setup of our experiment involving one risky asset which is predictable by the dividend yield and a risk-less asset. In Sect. 4, we compare the two methods and show that with short sale constraints, iterating on portfolio weights leads to numerically more accurate results.

2 The BGSS method

BGSS (2005) propose a simulation-based dynamic programming method for solving discrete-time multi-period portfolio choice problems involving non-standard preferences, a large number of assets with arbitrary return distribution, and, most importantly, a large number of state variables with potentially path-dependent or non-stationary dynamics. Their method does not rely on standard return distributions to generate closed-form expectations or on standard preferences to rule out path dependencies on wealth. Their method can

also be extended to accommodate Value-at-Risk and/or expected loss constraints, as done by van Binsbergen and Brandt (2006) who solve the asset allocation problem of an investment manager of a defined benefits pension plan who wishes to maximize the ratio of assets to liabilities while operating under regulatory risk constraints.

The BGSS method, as applied by van Binsbergen and Brandt (2006), essentially consists of five steps.

Step 1 Simulate a large number N of hypothetical sample paths with length T of the asset returns and the state variables that drive asset returns. These paths are simulated from the known, estimated, or bootstrapped joint dynamics of the returns and state variables, for example given by a Vector Auto Regression (VAR).

In period $T - 1$ the investor needs to decide on his portfolio weights for investment period T . These portfolio weights are based on the portfolio return moments for period T such as the mean, the variance, and the skewness. The return moments are dependent on the state variables in period $T - 1$, so we need to evaluate the conditional return moments.

Step 2a For a given set of portfolio weights, determine the conditional portfolio return moments for period T by simple OLS regressions. For example, to determine the conditional mean portfolio return for period T , we take the N simulated portfolio returns in period T and regress them on a constant and the simulated state variables in period $T - 1$. The fitted value of this regression gives us the conditional mean in period T as a linear function of the state variables in period $T - 1$. Similarly, if we want to determine the conditional portfolio variance for period T , we take the N simulated portfolio returns for period T , demean them with the conditional mean, square them, and regress them on a constant and the state variables in period $T - 1$. We can proceed with this method to obtain expressions for all relevant return moments.

Obviously, modeling the return moments as linear functions of the state variables may be too restrictive. This modeling assumption is easy to relax by including higher powers of the state variables on the right-hand side of the regression. For example, we can easily include the squared state variables or the cubed state variables.

Step 3a Approximate the utility function of the investor by a Taylor series expansion and take expectations. This results in an expression for the approximate expected utility value that involves the conditional return moments. Now substitute the functional expressions for the conditional return moments from step 2a to express the expected utility of the agent as a function of the state variables. This function then allows us to compute the expected utility in each simulation path.

We can also do steps 2a and 3a at once. So as an alternative to steps 2a and 3a we can implement the following step.

Step 2b and 3b Use the N simulated portfolio returns to directly compute the N utility values that these portfolio returns will result in. When we regress these utilities on a constant and the simulated state variables, we directly obtain the expected utility of the agent as a function of the state variables.

In steps 2 and 3, we have used portfolio returns, as opposed to single asset returns, in the regressions. To compute these portfolio returns we first had to pick a set of portfolio weights. However, as we are in the process of optimizing over these portfolio weights, we obviously do not know their optimal value yet. **Step 4** Repeat steps 2 and 3 for a grid of portfolio weights and pick for each path i those weights that maximize the expected utility of the investor.

Note that this grid-search is very robust and avoids a number of numerical issues about convergence that could emerge when taking first order conditions and iterating to a solution, as suggested in the original BGSS method. For a sufficiently fine grid, our grid search ensures finding a solution close to the global optimum.

Step 5 The algorithm then proceeds recursively backward until time zero. So we repeat the analysis for period $T - 2, T - 3, \dots, 0$.

Note that there are two ways to proceed with the backward recursion. When determining the optimal portfolio weights at time t we can either use the optimal *actions* of the agent from $t + 1$ through time $T - 1$ or equivalently, we can use the *expected utility* derived from these optimal actions. Using the optimized portfolio weights in the recursion is called portfolio weight iteration, whereas using the recursively estimated value function is called value function iteration. Both approaches are explained in more detail below.

An important question that has not been addressed is whether iterating on optimized portfolio weights leads to more or less accurate results than iterating on the value function. BGSS conjecture that iterating on the portfolio weights is superior but do not corroborate this claim through a comparative analysis. This article addresses this question. We compare the performance of both dynamic programming methods and show that iterating on portfolio weights leads to more accurate results, specifically when these portfolio weights are constrained by short sale restrictions. Value function iterations result in a lower variance but disproportionately higher bias of the solution, especially when risk aversion is high and the investment horizon is long. We note that the results we derive here, which favor iterating on portfolio weights and therefore confirm BGSS's claim, may be different for other dynamic programming methods, although the intuition underlying the superiority of iterating on optimized portfolio weight is fairly general.

3 Portfolio choice problem

In specifying the problem we use the same setup as in BGSS. Consider the portfolio choice at time t of an investor who maximizes the expected utility of wealth at some terminal date T by trading N risky assets and a risk-free asset (cash) at times $t, t + 1, \dots, T - 1$. The investor's problem can be written as:

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

subject to the sequence of budget constraints:

$$W_{s+1} = W_s \left(x'_s R_{s+1}^e + R^f \right) \quad \forall s \geq t, \tag{2}$$

where x_s is a vector of portfolio weights on the risky assets chosen at time s , R_{s+1}^e is the vector of excess returns on the N risky assets from time s to $s + 1$, and R^f is the gross return on the constant risk-free asset. The function $u(\cdot)$ represents the utility function of the investor over terminal wealth. The function $V_t(W_t, Z_t)$ is the value function and represents the expectation at *time* t , conditional on the state variable Z_t , of the utility of terminal wealth W_T generated by the current wealth W_t and the subsequent *optimal* portfolio weights.

It is relatively straightforward to write the multi-period problem as a single-period problem with utility $V_{t+1}(W_{t+1}, Z_{t+1})$ of next period's wealth W_{t+1} and information Z_{t+1} :

$$\begin{aligned} V_t(W_t, Z_t) &= \max_{\{x_s\}_{s=t}^{T-1}} E_t[u(W_T)] = \max_{x_t} E_t \left\{ \max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1}[u(W_T)] \right\} \\ &= \max_{x_t} E_t \left\{ V_{t+1} \left[W_t \left(x'_t R_{t+1}^e + R^f \right), Z_{t+1} \right] \right\}. \end{aligned} \tag{3}$$

This equation is the so-called Bellman equation. The first-order conditions (FOC) for an optimum are given by:

$$E_t \left\{ \partial_1 V_{t+1} \left[W_t \left(x'_t R_{t+1}^e + R^f \right), Z_{t+1} \right] R_{t+1}^e \right\} = 0. \tag{4}$$

These FOCs make up a system of non-linear equations that can be solved numerically for x_t . In the empirical example that we present, we will assume power utility (CRRA) over terminal wealth. In this case, the Bellman equation simplifies to:

$$\begin{aligned} V_t(W_t, Z_t) &= \max_{x_t} E_t \left\{ \max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1} \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \right\} \\ &= \max_{x_t} E_t \left\{ \max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1} \left[\frac{\left[W_{ts=t}^{T-1} \left(x'_s R_{s+1}^e + R^f \right) \right]^{1-\gamma}}{1-\gamma} \right] \right\} \tag{5} \\ &= \max_{x_t} E_t \left\{ \left[W_t \left(x'_t R_{t+1}^e + R^f \right) \right]^{1-\gamma} \max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1} \left[\frac{\left[\pi_{s=t+1}^{T-1} \left(x'_s R_{s+1}^e + R^f \right) \right]^{1-\gamma}}{1-\gamma} \right] \right\} \\ &= \max_{x_t} E_t \left\{ u(W_{t+1}) \psi_{t+1}(Z_{t+1}) \right\}, \end{aligned}$$

where

$$\frac{1}{1-\gamma} \psi_{t+1}(Z_{t+1}) \equiv \max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1} \left[\frac{\left[\pi_{s=t+1}^{T-1} (x'_s R_{s+1}^e + R^f) \right]^{1-\gamma}}{1-\gamma} \right]. \quad (6)$$

With power utility, the value function next period $V_{t+1}(W_{t+1}, Z_{t+1})$ can be expressed as the product of the utility of wealth $u(W_{t+1})$ and a function of the state variables $\psi_{t+1}(Z_{t+1})$. Further, as the CRRA utility function is homothetic in wealth, we can normalize $W_t = 1$. This simplifies the Bellman equation to:

$$\frac{1}{1-\gamma} \psi_t(Z_t) = \max_{x_t} E_t \left[u(x'_t R_{t+1}^e + R^f) \psi_{t+1}(Z_{t+1}) \right]. \quad (7)$$

The conditional expectations in each step of the backward induction are determined through regressions on a polynomial expansion of the state variables across simulation paths. Usually a first- or second- order Taylor expansion of the state variables suffices but higher order terms can be included to fit better the shape of the policy function at the center but at the cost of larger outliers at the boundaries. We use a first- and second-order Taylor expansion.

We can now illustrate the difference between iterating on portfolio weights and iterating on the value function. In period $T - 1$ we solve:

$$\max_{x_{T-1}} E_{T-1} \left[u(x'_{T-1} R_T^e + R^f) \right], \quad (8)$$

where, as explained above, the conditional expectation at time $T - 1$ of $u(x'_{T-1} R_T^e + R^f)$ is evaluated by a regression across simulation paths of the realized utility values on a set of basis functions of the state variables (steps 2b and 3b). This procedure leads to the optimal solution x_{T-1}^* for each simulation path and the corresponding values of the value function are given by:

$$\frac{1}{1-\gamma} \psi_{T-1}(Z_{T-1}) = E_{T-1} \left[u(x'_{T-1} R_T^e + R^f) \right]. \quad (9)$$

When we iterate on portfolio weights, on one hand, we store for each simulation path the optimal portfolio weight x_{T-1}^* and then solve in period $T - 2$:

$$\max_{x_{T-2}} E_{T-2} \left[\frac{\left[(x'_{T-2} R_{T-1}^e + R^f) (x'_{T-1} R_T^e + R^f) \right]^{1-\gamma}}{1-\gamma} \right]. \quad (10)$$

When we iterate on the value function, on the other hand, we store for each simulation path the value function $\psi_{T-1}(Z_{T-1})$ and then solve in period $T - 2$:

$$\max_{x_{T-2}} E_{T-2} \left[u \left(x'_{T-2} R_{T-1}^e + R^f \right) \psi_{T-1} (Z_{T-1}) \right]. \tag{11}$$

We can easily rewrite expression (11) as:

$$\begin{aligned} & \max_{x_{T-2}} E_{T-2} \left[u \left(x'_{T-2} R_{T-1}^e + R^f \right) \psi_{T-1} (Z_{T-1}) \right] \\ &= \max_{x_{T-2}} E_{T-2} \left[\left(x'_{T-2} R_{T-1}^e + R^f \right)^{1-\gamma} \hat{E}_{T-1} \left[u \left(x'_{T-1} R_T^e + R^f \right) \right] \right] \\ &= \max_{x_{T-2}} E_{T-2} \left[\left(x'_{T-2} R_{T-1}^e + R^f \right)^{1-\gamma} \hat{E}_{T-1} \left(\frac{\left(x'_{T-1} R_T^e + R^f \right)^{1-\gamma}}{1-\gamma} \right) \right] \end{aligned} \tag{12}$$

From expression (12) we can see that by the law of iterated expectations, in theory the two approaches are identical. In other words, when we are able to estimate the conditional moments exactly, it does not matter whether we iterate on portfolio weights or on the value function. However, in practice we approximate the conditional moments through regressions across simulation paths on a polynomial basis of the state variables. The main question we address here is whether the propagation of errors is different for these two methods.

A straightforward potential explanation for why it is better to iterate on portfolio weights when these weights are bounded by, for example, short sale constraints is the following. Suppose that in path i there is an error in the value function at time $T - 1$. Note that the value function is itself a conditional expectation, so an error in the value function means that this conditional expectation is not evaluated accurately when projecting realized utilities on the polynomial basis of the state variables. In other words, the polynomial expansion of the state variables is just an approximation of the true functional relation between the state variables and the conditional moments. This approximation leads to errors in the value function. As a consequence, the chosen portfolio weight in path i will also contain an error, and given the flat maximum of the utility function, this error will be large. However, when there are short sale constraints, the region of feasible portfolio weights is bounded and therefore the error is bounded as well. When we iterate on portfolio weights we therefore bound the accumulation of errors from period to period. On the other hand, when we iterate on the value function we are not using this convenient bound on the error and the accumulation of errors will be more pronounced.

4 Empirical example

We further explore the advantage of iterating on portfolio weights in the following example as presented in BGSS. We assume a restricted VAR as the quarterly data generating process:

$$\begin{bmatrix} r_{t+1}^c \\ d_{t+1} - p_{t+1} \end{bmatrix} = \begin{bmatrix} 0.227 & (0.95) \\ -0.155 & (-0.79) \end{bmatrix} + \begin{bmatrix} 0.060 & (0.87) \\ 0.958 & (17.02) \end{bmatrix} (d_t - p_t) + \begin{bmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+2} \end{bmatrix}$$

$$\begin{bmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+2} \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0.0060 & -0.0051 \\ -0.0051 & 0.0049 \end{bmatrix} \right), \tag{13}$$

where r_t^c is the log excess return on the value weighted CRSP index and $d_t - p_t$ denotes the log dividend yield of the index, computed from the sum of the past 12 monthly dividends and the current level of the index. The log dividend yield is the sole state variable (Z_t) in this example which, at time 0, we set equal to its unconditional mean. The VAR is estimated using quarterly data from January 1986 to December 1995. In parentheses are Newey and West (1987) adjusted t-statistics.

Tables 1 and 2 show the portfolio weights for values of risk aversion γ equal to 5, 10, 15, and 20, for an investment horizon T of 2, 4, 8, 12, 20, and 40 quarters and for the number of simulation paths N equal to 100, 1,000, and 10,000. Table 1 presents the results for the portfolio weight iteration method. Table 2 presents the results for the value function iteration methods. Tables 3 and 4 show the corresponding certainty equivalents. We use a first order polynomial expansion (linear) of the state variables in the regressions and find the optimal portfolio weight by a grid search over the interval [0,1] in steps of 0.01. A quadratic polynomial expansion leads to similar results. We run each simulation 80 times and report the average portfolio weight and certainty equivalent and their respective standard deviation over the 80 simulations.

In both tables, we also report the portfolio weights derived through the discretized state space method (DSS). We discretize the log dividend price ratio between -4.5 and -1.7 in fine steps of 0.1 and run for each point in the grid a simulation of 100,000 paths. We search for the optimal portfolio weight by a grid search over the interval [0,1] in steps of 0.01. We use this discretized state space method as the benchmark to compare portfolio weight iteration with value function iteration.

The tables make clear that for low values of the investment horizon T and risk aversion γ the two methods are almost equivalent and are reasonably accurate for N equal to 10,000. When the investment horizon and/or the risk aversion increase, both methods lose in precision. However, portfolio weight iteration seems to perform better than value function iteration. Specifically, it turns out that iterating on the value function leads to biased portfolio weights. To illustrate

Table 1 Portfolio weight iteration

		Weight in stocks			DSS
	γ	$N = 100$	$N = 1,000$	$N = 10,000$	
$T = 2$	5	0.6720(0.2401)	0.6750(0.0897)	0.6747(0.0255)	0.64
	10	0.3499(0.1614)	0.3430(0.0480)	0.3392(0.0138)	0.31
	15	0.2354(0.1062)	0.2316(0.0294)	0.2275(0.0088)	0.21
	20	0.1550(0.0694)	0.1714(0.0281)	0.1695(0.0081)	0.16
$T = 4$	5	0.7099(0.2304)	0.7175(0.1052)	0.7300(0.0284)	0.70
	10	0.3891(0.1653)	0.3719(0.0490)	0.3725(0.0158)	0.34
	15	0.2433(0.1055)	0.2503(0.0345)	0.2471(0.0107)	0.24
	20	0.1876(0.0952)	0.1841(0.0291)	0.1868(0.0087)	0.17
$T = 8$	5	0.7840(0.2349)	0.8408(0.1067)	0.8363(0.0348)	0.81
	10	0.4890(0.2129)	0.4334(0.0575)	0.4407(0.0212)	0.42
	15	0.3069(0.1564)	0.2939(0.0463)	0.2919(0.0165)	0.28
	20	0.2455(0.1161)	0.2155(0.0459)	0.2143(0.0240)	0.21
$T = 12$	5	0.8268(0.2159)	0.9189(0.0805)	0.9292(0.0328)	0.88
	10	0.5535(0.2037)	0.5208(0.0817)	0.5080(0.0327)	0.48
	15	0.4191(0.1990)	0.3560(0.0717)	0.3360(0.0363)	0.32
	20	0.2851(0.1463)	0.2568(0.0610)	0.2310(0.0446)	0.25
$T = 20$	5	0.9026(0.1598)	0.9688(0.0578)	0.9997(0.0022)	1.00
	10	0.6918(0.2140)	0.6623(0.1102)	0.6561(0.0531)	0.60
	15	0.5461(0.2233)	0.4917(0.1134)	0.4385(0.0671)	0.41
	20	0.4288(0.1394)	0.3279(0.1010)	0.2876(0.0731)	0.32
$T = 40$	5	0.9518(0.1009)	0.9999(0.0011)	1.0000(0.0000)	1.00
	10	0.8301(0.1625)	0.9172(0.0983)	0.9430(0.0477)	0.84
	15	0.7436(0.1953)	0.7915(0.0907)	0.7340(0.0915)	0.68
	20	0.6060(0.2106)	0.6421(0.1116)	0.5596(0.1450)	0.52

Portfolio weight in stocks at time 0 for a portfolio choice problem with a risk-free asset and one risky asset (stocks) with a predictable equity risk premium driven by the dividend yield as given in the text, 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 number of simulations N of 100, 1,000, and 10,000. The table reports averages and standard deviations (between brackets) of 80 simulations

Table 2 Value function iteration

		Weight in stocks			DSS
	γ	$N = 100$	$N = 1,000$	$N = 10,000$	
$T = 2$	5	0.6696(0.2376)	0.6754(0.0878)	0.6746(0.0256)	0.64
	10	0.3496(0.1625)	0.3450(0.0489)	0.3396(0.0133)	0.31
	15	0.2429(0.1088)	0.2326(0.0302)	0.2276(0.0083)	0.21
	20	0.1781(0.1373)	0.1726(0.0259)	0.1699(0.0082)	0.16
$T = 4$	5	0.7108(0.2250)	0.7204(0.1007)	0.7282(0.0264)	0.70
	10	0.3891(0.1500)	0.3713(0.0445)	0.3733(0.0153)	0.34
	15	0.2726(0.1420)	0.2488(0.0345)	0.2464(0.0093)	0.24
	20	0.3140(0.2938)	0.1893(0.0269)	0.1870(0.0082)	0.17
$T = 8$	5	0.7875(0.2194)	0.8403(0.0929)	0.8330(0.0313)	0.81
	10	0.4941(0.1648)	0.4426(0.0459)	0.4380(0.0152)	0.42
	15	0.4673(0.3016)	0.3059(0.0330)	0.2945(0.0108)	0.28
	20	0.6206(0.3797)	0.3681(0.2781)	0.2249(0.0084)	0.21

Table 2 continued

		Weight in stocks			DSS
	γ	$N = 100$	$N = 1,000$	$N = 10,000$	
$T = 12$	5	0.8385(0.2182)	0.9181(0.0786)	0.9339(0.0298)	0.88
	10	0.6318(0.1941)	0.5445(0.0576)	0.5328(0.0185)	0.48
	15	0.8200(0.2732)	0.5293(0.2302)	0.3909(0.0243)	0.32
	20	0.8478(0.3074)	0.8415(0.2886)	0.8825(0.2381)	0.25
$T = 20$	5	0.9293(0.1331)	0.9970(0.0125)	1.0000(0.0000)	1.00
	10	0.9596(0.1109)	0.9999(0.0011)	1.0000(0.0000)	0.60
	15	0.8455(0.3515)	1.0000(0.0000)	1.0000(0.0000)	0.41
	20	0.5000(0.5032)	0.6250(0.4872)	0.8875(0.3180)	0.32
$T = 40$	5	1.0000(0.0000)	1.0000(0.0000)	1.0000(0.0000)	1.00
	10	0.6125(0.4903)	0.7125(0.4555)	0.9750(0.1571)	0.84
	15	0.0750(0.2651)	0.0000(0.0000)	0.0000(0.0000)	0.68
	20	0.0125(0.1118)	0.0000(0.0000)	0.0000(0.0000)	0.52

Portfolio weight in stocks at time 0 for a portfolio choice problem with a risk-free asset and one risky asset (stocks) with a predictable equity risk premium driven by the dividend yield as given in the text, 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 number of simulations N of 100, 1,000, and 10,000. The table reports averages and standard deviations (between brackets) of 80 simulations

Table 3 Portfolio weight iteration

		Certainty equivalent		
	γ	$N = 100$ (%)	$N = 1,000$ (%)	$N = 10,000$ (%)
$T = 2$	5	8.99 (1.37)	8.80 (0.52)	8.70 (0.15)
	10	7.77 (0.97)	7.38 (0.25)	7.34 (0.07)
	15	7.12 (0.68)	6.94 (0.17)	6.91 (0.05)
	20	6.71 (0.50)	6.69 (0.15)	6.68 (0.04)
$T = 4$	5	9.76 (1.28)	8.93 (0.45)	8.89 (0.13)
	10	7.99 (0.73)	7.52 (0.22)	7.47 (0.05)
	15	7.27 (0.55)	7.01 (0.14)	6.97 (0.04)
	20	6.86 (0.56)	6.74 (0.11)	6.74 (0.04)
$T = 8$	5	9.97 (0.91)	9.28 (0.32)	9.19 (0.09)
	10	8.26 (0.69)	7.70 (0.21)	7.67 (0.06)
	15	7.53 (0.51)	7.13 (0.13)	7.10 (0.06)
	20	7.03 (0.50)	6.82 (0.15)	6.81 (0.07)
$T = 12$	5	9.99 (0.74)	9.51 (0.22)	9.42 (0.07)
	10	8.46 (0.59)	7.93 (0.19)	7.86 (0.07)
	15	7.73 (0.53)	7.32 (0.18)	7.20 (0.08)
	20	7.26 (0.52)	6.86 (0.20)	6.81 (0.15)
$T = 20$	5	10.21 (0.49)	9.68 (0.20)	9.64 (0.05)
	10	8.82 (0.40)	8.32 (0.23)	8.20 (0.09)
	15	8.08 (0.50)	7.65 (0.23)	7.42 (0.13)
	20	7.59 (0.42)	7.09 (0.24)	6.90 (0.16)

Table 3 continued

		Certainty equivalent		
γ		$N = 100$ (%)	$N = 1,000$ (%)	$N = 10,000$ (%)
$T = 40$	5	10.00 (0.26)	9.67 (0.09)	9.63 (0.03)
	10	9.00 (0.28)	8.72 (0.13)	8.64 (0.06)
	15	8.41 (0.28)	8.13 (0.16)	7.95 (0.15)
	20	7.85 (0.33)	7.63 (0.18)	7.36 (0.20)

Annualized certainty equivalent for a portfolio choice problem with a risk-free asset and one risky asset (stocks) with a predictable equity risk premium driven by the dividend yield as described in the text, 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 number of simulations N of 100, 1,000, and 10,000. The table reports averages and standard deviations (between brackets) of 80 simulations

Table 4 Value function iteration

		Certainty equivalent		
γ		$N = 100$ (%)	$N = 1,000$ (%)	$N = 10,000$ (%)
$T = 2$	5	8.99(1.37)	8.80(0.52)	8.70(0.15)
	10	7.75(0.96)	7.38(0.25)	7.34(0.07)
	15	7.10(0.72)	6.94(0.17)	6.91(0.05)
	20	6.48(1.52)	6.69(0.15)	6.68(0.04)
$T = 4$	5	9.73(1.28)	8.93(0.45)	8.89(0.13)
	10	7.95(0.71)	7.52(0.22)	7.47(0.05)
	15	7.14(0.60)	7.00(0.14)	6.97(0.04)
	20	5.90(2.47)	6.73(0.12)	6.73(0.04)
$T = 8$	5	9.90(0.91)	9.27(0.32)	9.19(0.09)
	10	8.10(0.70)	7.66(0.23)	7.67(0.06)
	15	6.60(1.80)	7.10(0.17)	7.09(0.10)
	20	4.13(3.54)	5.71(2.35)	6.78(0.15)
$T = 12$	5	9.89(0.74)	9.49(0.23)	9.42(0.07)
	10	8.11(0.77)	7.84(0.25)	7.85(0.09)
	15	6.10(1.80)	6.77(1.19)	7.06(0.28)
	20	3.87(2.69)	3.73(2.48)	3.39(2.06)
$T = 20$	5	10.10(0.50)	9.67(0.20)	9.64(0.05)
	10	8.36(0.69)	7.97(0.50)	7.93(0.22)
	15	6.01(1.59)	5.88(1.15)	5.39(0.74)
	20	3.71(2.06)	3.14(1.41)	2.40(1.29)
$T = 40$	5	9.89(0.29)	9.67(0.09)	9.65(0.03)
	10	7.34(0.87)	7.99(0.46)	8.10(0.16)
	15	4.70(1.12)	4.74(0.66)	4.22(0.78)
	20	3.37(1.14)	2.70(0.95)	2.00(0.73)

Annualized certainty equivalent for a portfolio choice problem with a risk-free asset and one risky asset (stocks) with a predictable equity risk premium driven by the dividend yield as given in the text, 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 number of simulations N of 100, 1,000, and 10,000. The table reports averages and standard deviations (between brackets) of 80 simulations

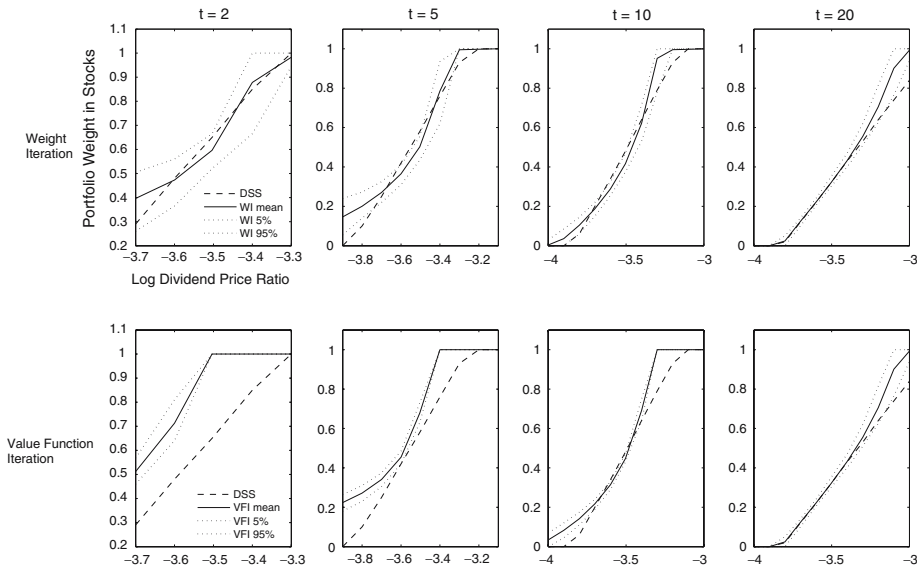


Fig. 1 Policy functions: we compare the policy functions of portfolio weight iteration (*first row*) and value function iteration (*second row*) for risk aversion $\gamma = 10$ and investment horizon $T = 20$ at $t = 2, t = 5, t = 10,$ and $t = 20$. We report the mean policy function (over 80 simulations) and the 5 and 95 percentiles. Note that the policy function at time t is in fact chosen at time $t - 1$. The policy function at $t = 2$ is therefore the policy function chosen after one period has passed. DSS is the policy function derived through the discretized state space approach

To illustrate this, Fig. 1 draws for investment horizon $T = 20$ and risk aversion $\gamma = 10$ the average optimal portfolio weight as a function of the log dividend price ratio for $t = 20, t = 10, t = 5,$ and $t = 2$ for both methods.¹ The average is taken over the 80 simulations. Further, the graphs shows the five and 95 percentiles of the policy functions. Finally, as a benchmark, the graphs depict the policy function that follows from the DSS approach. The figure clearly indicates that even though the variance around the average policy function seems to be smaller for the value function iteration, this policy function is also biased.

5 Conclusion

In this paper, we compare the numerical performance of value function iterations with portfolio weight iterations in the context of the simulation-based dynamic programming approach developed by BGSS (2005). We conclude that there is an advantage to using portfolio weight iterations, especially when risk aversion is high and the investment horizon is long. Value function iterations result in a lower variance but disproportionately higher bias of the solution.

¹ Note that the policy function for $t = 20$ is in fact chosen at time $t = 19$, but applies to the returns generated in period $t = 20$.

Acknowledgments We thank Amit Goyal, Ralph Kojien and an anonymous referee for helpful comments.

References

- Balduzzi, P., & Lynch, A. W. (1999). Transaction costs and predictability: some utility cost calculations. *Journal of Financial Economics*, *52*, 47–78.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance*, *55*, 225–264.
- van Binsbergen, J. H., & Brandt, M. W. (2006) *Optimal asset allocation in asset liability management*. Working Paper, Duke University.
- Brandt, M. W. (2005). Portfolio choice problems. In Y. Ait-Sahalia, & L. P. Hansen, (Eds.), *Handbook of Financial Econometrics*, forthcoming.
- 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. *Review of Financial Studies*, *18*, 831–873.
- Cochrane, J. H. (1989). The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives. *American Economic Review*, *79*, 319–337.
- Dammon, R. M., Spatt, C. S., & Zhang, H. H. (2001). Optimal consumption and investment with capital gains taxes. *Review of Financial Studies*, *14*, 583–616.
- Newey, W. K., & West K. D. (1987). A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*, 703–708.