

Numerical Solutions to Dynamic Portfolio Problems: The Case for Value Function Iteration using Taylor Approximation

Lorenzo Garlappi · Georgios Skoulakis

Accepted: 9 September 2008 / Published online: 25 September 2008
© Springer Science+Business Media, LLC. 2008

Abstract In a recent paper, van Binsbergen and Brandt (*Computational Economics*, 29, 355–367, 2007), using the method of Brandt et al. (*Review of Financial Studies*, 18, 831–873, 2005), argue, in the context of a portfolio choice problem with CRRA preferences, that value function iteration (VFI) is inferior to portfolio weight iteration (PWI), when a Taylor approximation is used. In particular, they report that the value function iteration produces highly inaccurate solutions when risk aversion is high and the investment horizon long. We argue that the reason for the deterioration of VFI is the high nonlinearity of the value function and illustrate that if one uses a natural and economically-motivated transformation of the value function, namely the *certainty equivalent*, the VFI approach produces very accurate results.

Keywords Portfolio choice · Numerical solution · Value function iteration

1 Introduction

Portfolio choice problems form a very important class of research problems in financial economics. Despite the great progress that has been made on the theoretical front of

L. Garlappi
Finance Department, McCombs School of Business, University of Texas at Austin, Austin,
TX 78712, USA
e-mail: lorenzo.garlappi@mcombs.utexas.edu

G. Skoulakis (✉)
Finance Department, Robert H. Smith School of Business, University of Maryland, College Park,
MD 20742, USA
e-mail: gskoulak@rhsmith.umd.edu

this literature,¹ a large number of interesting and relevant problems do not admit exact closed-form solutions and, therefore, call for analytical or numerical approximations. Indeed, a large number of papers have recently appeared that employ approximation techniques to tackle challenging portfolio choice problems. These approaches include, among others, numerical solution of a partial differential equation as in Brennan et al. (1997), log-linearization of the budget equation as in Campbell and Viceira (1999), state space discretization as in Balduzzi and Lynch (1999) and Barberis (2000), and, more recently, Malliavin calculus combined with Monte Carlo methods as in Detemple et al. (2003). A comprehensive survey of the recent portfolio choice literature can be found in Brandt (2004).

A common numerical approach to the solution of a portfolio choice problem is to use value function iteration (VFI) based on the associated Bellman equation. Typically, in each step of the iteration, one computes the (approximate) value function on a grid and then, using this information, the value function over the entire state space is approximated by interpolation or projection on a set of basis functions, such as high-order polynomials or radial basis functions. In contrast to this approach, Brandt et al. (2005) (BGSS hereafter) propose an alternative technique that relies on iteration of the (approximate) optimal portfolio weights. van Binsbergen and Brandt (2007) (vBB hereafter) further examine the performance of the BGSS approach in comparison to the traditional VFI approach, in the context of a portfolio choice problem with CRRA preferences. They find that the VFI approach that uses Taylor approximation produces less accurate results than the BGSS approach that uses portfolio weight iteration, especially for high levels of risk aversion and long investment horizons. In this paper, we argue that this conclusion applies to the specific way that vBB implement the VFI approach and the main reason for its poor performance is that they approximate the highly nonlinear value function using low-order polynomials. To mitigate this problem, we suggest using VFI on the *certainty equivalent* instead of the original value function. The certainty equivalent is a much less nonlinear function than the original value function, especially for high levels of risk aversion. In addition, the original value function takes values very close to zero for certain values of the state variable, especially for long horizons. In contrast, the certainty equivalent function is stable over time and never gets close to zero. Because of these two features, it is much easier to approximate the certainty equivalent using a set of basis functions, such as polynomials. To support our argument, we solve the same problem as in vBB using the state variable decomposition (SVD) approach developed in Garlappi and Skoulakis (2008), that performs VFI on the certainty equivalent. Our results support the conclusion that, when properly implemented, the VFI approach yields very accurate results *over the entire state space*.² We further point out that the Taylor approximation with respect to next-period wealth used by the BGSS approach is valid only when *future optimal portfolio allocations are independent of current wealth*.³ This implies that use of the BGSS approach is problematic when wealth is not a redundant state variable, as is the

¹ Prominent examples of this literature are the seminal papers by Merton (1973) and Cox and Huang (1989) that use dynamic programming and martingale techniques, respectively.

² This point is worth emphasizing as vBB report results *only* for the mean of the state variable.

³ For further discussion of this issue, see Garlappi and Skoulakis (2008).

case for CRRA preferences. This excludes a number of important applications including realistic problems with taxes, transaction costs, and labor income. In contrast, the Taylor approximation used by the SVD approach is valid even when future optimal portfolio allocations do depend of current wealth.

The rest of the paper proceeds as follows. In Sect. 2, we describe the portfolio choice problem for an investor with CRRA preferences facing an investment opportunity set that consists of one risk-free and one risky asset with predictable returns. In Sect. 3, we describe the SVD approach, while in Sect. 4, we report empirical results that illustrate the numerical accuracy of the value function iteration that uses the certainty equivalent, as applied by the SVD method. We briefly conclude in Sect. 5.

2 Portfolio Choice Problem with CRRA Preferences

Consider an investor whose objective is to maximize expected utility of wealth at some terminal date T .⁴ The investor’s preferences are described by a utility function U . The investment opportunity set consists of a risk-free asset and N risky assets. The gross interest rate on the risk-free asset is denoted by R_f , the vector of excess returns on the risky assets from time t to time $t + 1$ is denoted by R_{t+1} , and the investor’s wealth at time t is denoted by W_t . At time 0, the investor’s optimization problem can then be expressed as:

$$\max_{\{\omega_t\}_{t=0}^{T-1}} E_0 [U(W_T)] \tag{1}$$

where ω_t is the N -dimensional vector of portfolio weights on the risky assets at time t . The evolution of wealth is described by the budget equation

$$W_{t+1} = W_t(R_f + \omega_t' R_{t+1}). \tag{2}$$

We assume that the investment opportunity set varies over time and therefore, besides the wealth W_t , there are additional state variables that govern the evolution of returns over time. Denoting these state variables by S_t , we have that the value function, or indirect utility of wealth, denoted by J_t , depends on both W_t and S_t . The investor’s problem is then characterized by the sequence of problems

$$J_t(W_t, S_t) = \max_{\{\omega_\tau\}_{\tau=t}^{T-1}} E_t [U(W_T)] \tag{3}$$

subject to the budget Eq. 2. The problem is then equivalent to the Bellman equation

$$J_t(W_t, S_t) = \max_{\omega_t} E_t [J_{t+1}(W_{t+1}, S_{t+1})]. \tag{4}$$

⁴ Since the main purpose of the paper is to facilitate a comparison with the PWI approach of [van Binsbergen and Brandt \(2007\)](#), we adopt their framework and ignore intermediate consumption. See [Garlappi and Skoulakis \(2008\)](#) for a problem that incorporates consumption in the context of recursive preferences.

Typically, the solution to the Bellman equation relies on backward iteration on the value function J_t starting from the terminal condition $J_T(W_T, S_T) = U(W_T)$. However, it is worth pointing out that, instead of J_t , one can use *any* strictly monotonic transformation of J_t in the backward recursion; there is no loss of information in using such transformation. The advantage of using a transformation is that one can control the nonlinearity of the function used in the recursion. This is an important point because this function has to be approximated in each step of the backward induction and, therefore, avoiding hard-to-approximate highly nonlinear functions can be critical in achieving computational accuracy and efficiency. In the present context of portfolio choice, there is a natural candidate for such a transformation, namely the inverse of the utility function U^{-1} . Under this choice, it follows that the transformed value function $V_t = U^{-1}(J_t)$ is nothing but the certainty equivalent of J_t , implicitly defined by

$$J_t(W_t, S_t) = U(V_t(W_t, S_t)). \tag{5}$$

To facilitate a comparison with the results in vBB, we restrict our analysis to the case of CRRA preferences. We assume that the investor has power utility function with coefficient of relative risk aversion γ :

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}. \tag{6}$$

It is well-known that the homotheticity of CRRA preferences implies that the value function J_t is separable in its two arguments, W_t and S_t . To see this, let us guess that the value function J_t is of the form

$$J_t(W_t, S_t) = \frac{W_t^{1-\gamma}}{1-\gamma} \mathcal{V}_t(S_t)^{1-\gamma} \tag{7}$$

or, equivalently, the certainty equivalent is linear in wealth:

$$V_t(W_t, S_t) = W_t \mathcal{V}_t(S_t). \tag{8}$$

The function $\mathcal{V}_t(S_t)$, that captures the dependence of the certainty equivalent function $V_t(W_t, S_t)$ on the state variable S_t , is referred to as the *reduced* certainty equivalent function. Substituting the expression (8) for $V_t(W_t, S_t)$ in Eq. 4 and using the budget Eq. 2, we obtain

$$\frac{W_t^{1-\gamma}}{1-\gamma} \mathcal{V}_t(S_t)^{1-\gamma} = \max_{\omega_t} E_t \left[\frac{[W_t(R_f + \omega'_t R_{t+1})]^{1-\gamma}}{1-\gamma} \mathcal{V}_{t+1}(S_{t+1})^{1-\gamma} \right].$$

Then, wealth W_t cancels and we obtain the reduced Bellman equation involving \mathcal{V}_t

$$\frac{1}{1-\gamma} \mathcal{V}_t(S_t)^{1-\gamma} = \max_{\omega_t} E_t \left[(R_f + \omega'_t R_{t+1})^{1-\gamma} \frac{1}{1-\gamma} \mathcal{V}_{t+1}(S_{t+1})^{1-\gamma} \right]. \tag{9}$$

Therefore, the value function J_t indeed has the separable form (7) where the reduced certainty equivalent function \mathcal{V}_t solves the Eq. 9. In particular, this fact implies that optimal portfolio weights are *independent* of wealth (current or past). This effectively eliminates wealth from the menu of state variables. Consistent with (7), the terminal condition on the reduced certainty equivalent function is $\mathcal{V}_T(S_T) = 1$.

In contrast, vBB do not use the certainty equivalent transformation and state the Bellman equation as⁵

$$\mathcal{J}_t(S_t) = \max_{\omega_t} E_t \left[(R_f + \omega'_t R_{t+1})^{1-\gamma} \mathcal{J}_t(S_{t+1}) \right] \tag{10}$$

where \mathcal{J}_t is the reduced value function implicitly defined by

$$J_t(W_t, S_t) = W_t^{1-\gamma} \mathcal{J}_t(S_t). \tag{11}$$

Note that the functions \mathcal{V}_t and \mathcal{J}_t are related through the equation

$$\mathcal{J}_t(S_t) = \frac{1}{1-\gamma} \mathcal{V}_t(S_t)^{1-\gamma}. \tag{12}$$

As suggested by the presence of the power $1 - \gamma$ in the preceding equation, it turns out that \mathcal{J}_t is much more nonlinear—and thus much harder to approximate—than \mathcal{V}_t , especially for high values of the risk aversion coefficient γ . We argue that this is the main reason behind the inferior performance of the value function iteration as implemented in vBB. While the two versions of the Bellman equation given by (9) and (10) are mathematically equivalent, working with Eq. 9 is preferable since \mathcal{V}_t is much easier to approximate. As illustrated in Eqs. 8 and 11, the homotheticity of CRRA preferences implies that the value function $J_t(W_t, S_t)$ and the certainty equivalent $V_t(W_t, S_t)$ are determined by their reduced counterparts $\mathcal{J}_t(S_t)$ and $\mathcal{V}_t(S_t)$, respectively. Hence, from this point on and with a slight abuse of terminology, we drop the qualifier “reduced” and simply refer to the functions \mathcal{J}_t and \mathcal{V}_t as the value function and the certainty equivalent function, respectively.

Although quite general, the methodology discussed in this paper will focus on the example with predictable risky asset returns analyzed in BGSS and vBB. Specifically, we assume that there is one risky asset and its log excess returns, denoted by r_t , are predictable by the log dividend yield, denoted by z_t . The joint dynamics of $(r_t, z_t)'$ are described by a restricted VAR(1) system as follows

$$\begin{bmatrix} r_{t+1} \\ z_{t+1} \end{bmatrix} = \begin{bmatrix} a_r & b_r \\ a_z & b_z \end{bmatrix} \begin{bmatrix} 1 \\ z_t \end{bmatrix} + \begin{bmatrix} \varepsilon_{r,t+1} \\ \varepsilon_{z,t+1} \end{bmatrix}, \quad \begin{bmatrix} \varepsilon_{r,t+1} \\ \varepsilon_{z,t+1} \end{bmatrix} \sim N \left(0, \begin{bmatrix} \sigma_r^2 & \sigma_{rz} \\ \sigma_{rz} & \sigma_z^2 \end{bmatrix} \right). \tag{13}$$

⁵ Strictly speaking, our Eq. 10 is a restatement of Eq. 7 in vBB which reads as $\frac{1}{1-\gamma} \psi_t(S_t) = \max_{\omega_t} E_t \left[\frac{1}{1-\gamma} (R_f + \omega'_t R_{t+1})^{1-\gamma} \psi_t(S_{t+1}) \right]$. Our function $\mathcal{J}_t(S_t)$ is related to the function $\psi_t(S_t)$ in vBB through $\mathcal{J}_t(S_t) = \frac{1}{1-\gamma} \psi_t(S_t)$. Working with \mathcal{J}_t facilitates better interpretation of our results, since \mathcal{J}_t , like \mathcal{V}_t , is an increasing function of the state variable S_t in the application that follows.

The above setup implies that the predictor z_t fully describes the investment opportunity set and therefore constitutes the only state variable. The log excess returns are defined by $r_t = \log(R_{g,t}) - \log(R_f)$ where $R_{g,t}$ denotes the gross risky asset return at time t . It follows that the excess return R_t is given by

$$R_t = R_{g,t} - R_f = R_f(e^{r_t} - 1). \tag{14}$$

To illustrate the perils of working directly with the value function, in Figure 1 we report the value function $\mathcal{J}_0(z)$ (left panels) and the certainty equivalent function $\mathcal{V}_0(z)$ (right panels) for three choices of the investment horizon $T = 20, 30,$ and 40 quarters, when preferences are CRRA and the risky asset returns are predictable by the log dividend yield z , which is the exogenous state variable. The data generating process and the parameters used are described in Eqs. 13 and 27, respectively, and the coefficient of relative risk aversion is set equal to $\gamma = 15$. The solution is based on backward




Fig. 1 Value function and certainty equivalent. This figure presents the value function $\mathcal{J}_0(z)$ and the certainty equivalent function $\mathcal{V}_0(z)$ (see Eqs. 10 and 12) for three choices of the investment horizon $T = 20, 30, 40$. The preferences are CRRA and the returns predictable by the log dividend yield z . The data generating process and the parameters used are described in Eqs. 13 and 27, respectively. The annualized risk-free rate is set equal to 6% and the coefficient of relative risk aversion is set equal to $\gamma = 15$. The solution is based on backward induction on the certainty equivalent. The expectations in the Bellman equation are computed using Gauss-Hermite quadrature with 6 nodes

induction on the certainty equivalent function $\mathcal{V}_t(z)$ using a quadrature-based discretized state space approach (see Sect. 4 for details). As the figure illustrates, the value function $\mathcal{J}_0(z)$ is a more nonlinear function of the state variable z compared to the certainty equivalent function $\mathcal{V}_0(z)$, making it harder to approximate. Furthermore, the value function assumes values very close to zero, especially for longer horizons. This is an undesirable property when an approximation is needed because the value function might end up being approximated by a function that takes on both positive and negative values. In contrast, the economically motivated transformation suggested by the certainty equivalent addresses both of these problems at once: the certainty equivalent function is easier to approximate and is far away from zero, reducing the risk of perverse approximations.

In the next section, we introduce the SVD approach of [Garlappi and Skoulakis \(2008\)](#) and in the following we implement it numerically for the case of CRRA preferences and predictable returns.

3 The State Variable Decomposition (SVD) Approach

The SVD approach uses value function iteration on the certainty equivalent \mathcal{V}_t based on the Bellman Eq. 9. The main purpose of SVD is to provide a computationally efficient and accurate method for computing the conditional expectations involved in (9). Specifically, SVD decomposes the random variables r_{t+1} and z_{t+1} (or equivalently the state variables W_{t+1} and z_{t+1}) into their conditional means and the associated shocks and uses a Taylor approximation to achieve a separation between quantities that depend *only* on choice variables and quantities that depend only on shocks to the state variables. As a result of this separation, the computation of the conditional expectations in (9) reduces to the computation of conditional expectations that depend only on shocks and not on choice variables. These expectations can then easily be computed, either analytically or numerically.

To illustrate the SVD approach in more detail, let us use the definition of U and Eq. 14 to rewrite the Bellman equation (9) as

$$U(\mathcal{V}_t(z_t)) = \max_{\omega_t} E_t \left[U \left(R_f(1 + \omega_t(e^{r_{t+1}} - 1)) \mathcal{V}_{t+1}(z_{t+1})^{1-\gamma} \right) \right]. \tag{15}$$

Suppose that an approximation to \mathcal{V}_{t+1} has already been obtained for the period $t + 1$. The goal is to obtain an approximation to \mathcal{V}_t using (15). We will implement two versions of the SVD method. The first, which we label *Partial SVD (PSVD)*, is based on the decomposition of only the return r_{t+1} in (15). The second, which we label *Full SVD (FSVD)*, is based on the decomposition of *both* the return r_{t+1} and the state variable z_{t+1} .

3.1 The Partial SVD (PSVD) Approach

The PSVD approach consists of the following three steps. First, we decompose the risky asset log excess return r_t into its conditional mean and the corresponding shock:

$$r_{t+1} = \mu_{r,t} + \varepsilon_{r,t+1} \tag{16}$$

where $\mu_{r,t} = E_t[r_{t+1}] = a_r + b_r z_t$ and the shock $\varepsilon_{r,t+1}$ satisfies $E_t[\varepsilon_{r,t+1}] = 0$, as it follows from the VAR system (13). In the second step, we separate the choice variable ω_t from the shock $\varepsilon_{r,t+1}$. Define the function

$$G_{t+1}(\varepsilon_{r,t+1}) = U(g_{t+1}(\varepsilon_{r,t+1})) \tag{17}$$

where

$$g_{t+1}(\varepsilon_{r,t+1}) = R_f(1 + \omega_t(e^{\mu_t + \varepsilon_{r,t+1}} - 1)) \tag{18}$$

and use the M -th order Taylor expansion of $G_{t+1}(\varepsilon_{r,t+1})$ with respect to $\varepsilon_{r,t+1}$ around 0 to obtain

$$G_{t+1}(\varepsilon_{r,t+1}) \approx \sum_{m=0}^M \frac{1}{m!} \frac{d^m G_{t+1}(0)}{d\varepsilon_{r,t+1}^m} \varepsilon_{r,t+1}^m. \tag{19}$$

The derivatives of the composite function G_{t+1} evaluated at 0 can be efficiently computed through a recursive scheme known as the Faá di Bruno formula.⁶ Note that the terms $\frac{d^m G_{t+1}(0)}{d\varepsilon_{r,t+1}^m}$ depend on the choice variable ω_t but *not* on the shock $\varepsilon_{r,t+1}$, and thus the desired separation has been achieved. Using the Bellman equation (15) and the approximation (19), we obtain the approximate Bellman equation

$$U(\mathcal{V}_t(z_t)) = \max_{\omega_t} \sum_{m=0}^M \frac{1}{m!} \frac{d^m G_{t+1}(0)}{d\varepsilon_{r,t+1}^m} E_t \left[\varepsilon_{r,t+1}^m \mathcal{V}_{t+1}(z_{t+1})^{1-\gamma} \right]. \tag{20}$$

The above step is important in terms of computational efficiency since the conditional expectations in (20) do not involve the choice variable ω_t . This means that the expectations have to be computed only once for each grid point in the state space, at time t , and do not have to be recomputed in each step of the optimization search.

The third, and final, step of the PSVD method addresses the computation of the conditional expectations $E_t \left[\varepsilon_{r,t+1}^m \mathcal{V}_{t+1}(z_{t+1})^{1-\gamma} \right]$ in the approximate Bellman equation (20). There are two alternatives for completing this step. The first possibility is to use quadrature methods, as in Judd (1998), to compute the conditional expectations in (20). According to (13), the shock vector $\varepsilon_{t+1} = (\varepsilon_{r,t+1}, \varepsilon_{z,t+1})$ is jointly normally distributed given the information at time t and therefore Gauss-Hermite quadrature can be used to compute the required expectations. A second possibility for computing the conditional expectations is to use the method of parameterized expectations of Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) that uses simulation and cross-sectional regressions. The same method is used by the BGSS approach. If

⁶ For details on the use of the Faá di Bruno formula, see Garlappi and Skoulakis (2008) and Savits (2006).

we use quadrature then the method is termed the PSVD-Q approach, while if we use simulation the method is termed the PSVD-S approach.⁷

3.2 The Full SVD (FSVD) Approach

In contrast to the PSVD approach, the first step of the FSVD approach decomposes both the return r_{t+1} and the state variable z_{t+1} into their expected and unexpected components, as follows

$$r_{t+1} = \mu_{r,t} + \varepsilon_{r,t+1}, \quad \mu_{r,t} = a_r + b_r z_t, \quad E_t[\varepsilon_{r,t+1}] = 0 \tag{21}$$

$$z_{t+1} = \mu_{z,t} + \varepsilon_{z,t+1}, \quad \mu_{z,t} = a_z + b_z z_t, \quad E_t[\varepsilon_{z,t+1}] = 0. \tag{22}$$

In the second step, we separate the choice variable ω_t from the shock $\varepsilon_{t+1} = (\varepsilon_{r,t+1}, \varepsilon_{z,t+1})$. To achieve this goal, we define the function

$$H_{t+1}(\varepsilon_{t+1}) \equiv U(h_{t+1}(\varepsilon_{t+1})), \tag{23}$$

where

$$h_{t+1}(\varepsilon_{t+1}) \equiv R_f (1 + \omega_t (e^{\mu_{r,t} + \varepsilon_{r,t+1}} - 1)) \mathcal{V}_{t+1}(\mu_{z,t} + \varepsilon_{z,t+1}). \tag{24}$$

The M -th order Taylor approximation of $H_{t+1}(\varepsilon_{t+1})$ with respect to ε_{t+1} around $(0, 0)$ is

$$H_{t+1}(\varepsilon_{t+1}) \approx \sum_{0 \leq m_1 + m_2 \leq M} \frac{1}{m_1! m_2!} H_{t+1}^{(m_1, m_2)}(0, 0) \varepsilon_{r,t+1}^{m_1} \varepsilon_{z,t+1}^{m_2}, \tag{25}$$

where $H_{t+1}^{(m_1, m_2)}(0, 0)$ denotes the (m_1, m_2) -partial derivative of $H_{t+1}(\varepsilon_{t+1})$ evaluated at $(0, 0)$. Using the Bellman equation (15) and the approximation (25), we obtain the approximate Bellman equation

$$U(\mathcal{V}_t(z_t)) = \max_{\omega_t} \sum_{0 \leq m_1 + m_2 \leq M} \frac{1}{m_1! m_2!} H_{t+1}^{(m_1, m_2)}(0, 0) E_t \left[\varepsilon_{r,t+1}^{m_1} \varepsilon_{z,t+1}^{m_2} \right]. \tag{26}$$

The partial derivatives of the composite function H_{t+1} are computed in an efficient fashion using the multidimensional version of the Faà di Bruno formula. Note that the derivatives $H_{t+1}^{(m_1, m_2)}(0, 0)$ depend on the portfolio weights ω_t but *not* on the shocks ε_{t+1} . The last step of the generic SVD approach concerns the computation of the conditional expectations $E_t \left[\varepsilon_{r,t+1}^{m_1} \varepsilon_{z,t+1}^{m_2} \right]$. It follows from (13) that, conditionally on

⁷ In our implementation, the PSVD-S method has resulted in inferior performance compared to FSVD and PSVD-Q. An additional advantage of the FSVD and PSVD-Q methods is that they are not subject to Monte-Carlo error. However, in cases in which analytical computations are not feasible, the PSVD approach is flexible enough to accommodate a simulation-based computation of expectations. For details see [Garlappi and Skoulakis \(2008\)](#).

the information available at time t , the vector $(\varepsilon_{r,t+1}, \varepsilon_{z,t+1})$ is normally distributed. Hence, the expectation in (26) is available in closed form and can be computed efficiently using the recursive algorithm in Willink (2005) and Savits (2006).

3.3 Solution of the Problem

The solution to the portfolio choice problem proceeds as follows. We start at time T using the terminal condition $\mathcal{V}_T(z_T) = 1$. At each time t , we select a set of grid points in the state space. Then following the three steps outlined above, we solve the approximate Bellman equation for each point on the grid to obtain the values of the function \mathcal{V}_t . Using these values, we obtain an approximation to the certainty equivalent function \mathcal{V}_t over the entire state space. This can be done by interpolation or projection on a set of basis functions such as polynomials or radial basis functions. The procedure is then repeated until we reach time 0.

While the SVD and BGSS approaches are both based on a Taylor expansion, there are some important differences between them. First, while the BGSS method uses an expansion based on a non-zero mean random quantity, specifically the risky portfolio return $\omega'_t R_{t+1}$, the SVD method uses an expansion based on the shock $\varepsilon_t = (\varepsilon_{r,t}, \varepsilon_{z,t})'$ which, by definition, has zero mean. Second, while the original BGSS method uses portfolio weight iteration and another version, as presented in vBB, uses value function iteration, the SVD method exclusively uses value function iteration. However, we should emphasize here that the version of BGSS that uses portfolio weight iteration is not valid in a generic portfolio choice problem. The reason is that the Taylor expansion used in BGSS⁸ is valid only when all future optimal portfolio weights are independent of current wealth. While this is the case for the standard portfolio choice problem with CRRA preferences and no frictions, in general, future optimal allocations *do* depend on current wealth. For example, this will be the case if preferences are CARA, if there are taxes and/or transaction costs, or if there is labor income. This dependence on wealth prevents the application of the BGSS portfolio weight iteration to generic portfolio choice problems.⁹ Third, while BGSS uses simulation to compute conditional expectations, SVD is more flexible allowing for simulation-based, quadrature-based, or fully analytical computation of conditional expectations. A main disadvantage of the BGSS approach is the rather large associated Monte Carlo error, as can be seen in Tables 2 and 4 in vBB.

4 Numerical Results

We illustrate the SVD approach using the setup described in Sect. 2 and the parameter values used in vBB. Specifically, we use a risk-free rate equal to 6% in annualized terms and the following VAR parameters

⁸ Specifically, see Eqs. 16 and 17 on p. 843 in Brandt et al. (2005).

⁹ See Garlappi and Skoulakis (2008) for a detailed discussion of this issue.

$$\begin{bmatrix} a_r & b_r \\ a_z & b_z \end{bmatrix} = \begin{bmatrix} 0.227 & 0.060 \\ -0.155 & 0.958 \end{bmatrix}, \quad \begin{bmatrix} \sigma_r^2 & \sigma_{rz} \\ \sigma_{rz} & \sigma_z^2 \end{bmatrix} = \begin{bmatrix} 0.0060 & -0.0051 \\ -0.0051 & 0.0049 \end{bmatrix}. \quad (27)$$

We impose no borrowing and short-sale constraints and therefore restrict all portfolio weights between 0 and 1. In what follows, we implement two versions of the SVD approach. The first, referred to as FSVD, is the full version and uses a full decomposition of both the return and the predictor and computes expectations analytically. The second, referred to as PSVD-Q, decomposes only the portfolio return and uses a quadrature-based method to numerically compute the resulting conditional expectations in (20). For comparison purposes, we also solve the problem using a Discretized State Space (DSS) approach that uses a quadrature method to directly compute expectations in the Bellman equation (15), as in Judd (1998), without resorting to Taylor approximations. We refer to this method as DSS-Q and use it as our benchmark. To achieve high accuracy, we use 500 grid points and linear interpolation of the certainty equivalent function $\mathcal{V}_t(z_t)$ when we implement the DSS-Q approach. The SVD approach uses a Taylor expansion of order $M = 4, 6, \text{ and } 8, 100$ grid points on the state space, and approximates the certainty equivalent function $\mathcal{V}_t(z_t)$ with a polynomial of degree 12 in z_t . All numerical computations of expectations use Gauss-Hermite quadrature with 6 nodes in each dimension.

We present the annualized certainty equivalent return (CER) and the optimal allocation to the risky asset obtained by all three methods. The CER is computed through extensive simulation. Specifically, for an investment horizon of T quarters and an initial predictor value equal to z_0 , we compute the estimated expected utility of terminal wealth, $\hat{E}_0[U(W_T(z_0))]$ by Monte Carlo simulation as follows:¹⁰

$$\hat{E}_0[U(W_T(z_0))] = \frac{1}{I} \sum_{i=1}^I U(W_T^i(z_0)), \quad (28)$$

$$W_T^i(z_0) = \prod_{t=0}^{T-1} \left(1 + \omega_t^*(z_t^i) R_{t+1}^i\right) \quad (29)$$

where I is a large integer, $\{(z_t^i, R_{t+1}^i) : t = 0, \dots, T - 1\}$, for each $i = 1, \dots, I$, is a simulated path of predictors and returns according to the data generating process (13) and the parameter values stated in (27) and $\omega_t^*(z_t^i)$ is the corresponding optimal portfolio. Each simulated path starts at $z_0^i = z_0$. To minimize the impact of simulation error on our results, we use a very large number of repetitions, $I = 10^7$. The CER is then computed as

$$CER_T(z_0) = \left(U^{-1} \left(\hat{E}_0[U(W_T(z_0))] \right) \right)^{\frac{4}{T}} - 1. \quad (30)$$

We consider four choices for the investment horizon, $T = 10, 20, 30, \text{ and } 40$ quarters and two levels of relative risk aversion, $\gamma = 5 \text{ and } 15$. Table 1 reports the annualized

¹⁰ Since the initial value of wealth, W_0 , does not affect the CER calculation, we simply assume $W_0 = 1$.

Table 1 Annualized certainty equivalent return

		$\gamma = 5$					$\gamma = 15$				
		z_{10}	z_{30}	z_{50}	z_{70}	z_{90}	z_{10}	z_{30}	z_{50}	z_{70}	z_{90}
<i>T</i> = 10											
DSS-Q		6.08	6.43	7.22	8.72	12.03	6.03	6.15	6.43	6.97	8.26
PSVD-Q	<i>M</i> = 4	6.08	6.43	7.22	8.72	12.04	6.03	6.15	6.43	6.97	8.26
	<i>M</i> = 6	6.08	6.43	7.22	8.72	12.04	6.03	6.15	6.43	6.97	8.26
	<i>M</i> = 8	6.08	6.43	7.22	8.71	12.04	6.03	6.15	6.43	6.97	8.25
FSVD	<i>M</i> = 4	6.08	6.43	7.22	8.72	12.04	6.03	6.15	6.43	6.97	8.26
	<i>M</i> = 6	6.08	6.43	7.23	8.72	12.04	6.03	6.15	6.43	6.97	8.26
	<i>M</i> = 8	6.08	6.43	7.22	8.71	12.04	6.03	6.15	6.43	6.97	8.25
<i>T</i> = 20											
DSS-Q		6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.35	8.76
PSVD-Q	<i>M</i> = 4	6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.35	8.74
	<i>M</i> = 6	6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.35	8.75
	<i>M</i> = 8	6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.34	8.77
FSVD	<i>M</i> = 4	6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.35	8.73
	<i>M</i> = 6	6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.35	8.75
	<i>M</i> = 8	6.34	6.92	7.84	9.29	12.09	6.12	6.35	6.72	7.34	8.77
<i>T</i> = 30											
DSS-Q		6.65	7.34	8.26	9.57	11.91	6.25	6.57	7.01	7.70	9.17
PSVD-Q	<i>M</i> = 4	6.65	7.34	8.26	9.57	11.91	6.25	6.56	7.01	7.69	9.09
	<i>M</i> = 6	6.65	7.34	8.26	9.57	11.91	6.26	6.56	7.01	7.68	9.13
	<i>M</i> = 8	6.65	7.34	8.26	9.57	11.91	6.26	6.57	7.01	7.70	9.18
FSVD	<i>M</i> = 4	6.65	7.34	8.26	9.57	11.92	6.26	6.56	7.02	7.69	9.18
	<i>M</i> = 6	6.65	7.34	8.26	9.57	11.91	6.26	6.56	7.01	7.70	9.17
	<i>M</i> = 8	6.65	7.34	8.26	9.57	11.91	6.26	6.57	7.01	7.71	9.18
<i>T</i> = 40											
DSS-Q		6.95	7.67	8.53	9.69	11.67	6.40	6.78	7.26	7.98	9.43
PSVD-Q	<i>M</i> = 4	6.95	7.67	8.53	9.69	11.67	6.40	6.78	7.25	7.93	9.27
	<i>M</i> = 6	6.95	7.67	8.53	9.69	11.67	6.40	6.78	7.27	7.98	9.38
	<i>M</i> = 8	6.95	7.67	8.53	9.69	11.67	6.40	6.77	7.26	7.98	9.42
FSVD	<i>M</i> = 4	6.95	7.67	8.53	9.69	11.67	6.41	6.78	7.26	7.99	9.45
	<i>M</i> = 6	6.95	7.67	8.53	9.69	11.67	6.41	6.78	7.27	7.99	9.41
	<i>M</i> = 8	6.95	7.67	8.53	9.69	11.67	6.40	6.78	7.25	7.98	9.43

The table reports the annualized certainty equivalent returns obtained by the quadrature-based discretized state space method (referred to as DSS-Q), the partial SVD approach that uses quadrature (referred to as PSVD-Q), and the full SVD approach (referred to as FSVD). The preferences are CRRA and the returns predictable by the log dividend yield. The data generating process and the parameters used are described in Eqs. 13 and 27, respectively. The annualized risk-free rate is set equal to 6%. We consider two values for the coefficient of relative risk aversion, $\gamma = 5$, and 15 and four values for the time horizon, $T = 10, 20, 30$, and 40. We present the results for 5 different initial values of the log dividend yield z . Each value corresponds to the p -th percentile of the unconditional distribution of z , denoted by z_p , where p takes values 10, 30, 50, 70, and 90.

CER and Table 2 reports the optimal allocations for 5 different initial values of the log dividend yield z . Each value corresponds to the p -th percentile of the unconditional distribution of z , where p takes values 10, 30, 50, 70, and 90.

As is evident from Tables 1 and 2, both versions of the SVD approach are very accurate across *all* investment horizons, levels of risk aversion, and levels of the initial value of the predictor. Specifically, we observe that the absolute difference in certainty

Table 2 Optimal allocation

		$\gamma = 5$					$\gamma = 15$				
		z_{10}	z_{30}	z_{50}	z_{70}	z_{90}	z_{10}	z_{30}	z_{50}	z_{70}	z_{90}
<i>T = 10</i>											
DSS-Q		0.0	13.2	42.8	72.5	100.0	0.0	4.5	15.6	27.1	44.5
PSVD-Q	<i>M</i> = 4	0.0	13.3	42.7	72.4	100.0	0.0	5.1	16.1	27.4	44.2
	<i>M</i> = 6	0.0	13.2	42.9	72.5	100.0	0.0	4.5	15.5	27.1	44.7
	<i>M</i> = 8	0.0	13.2	42.8	72.5	100.0	0.0	4.4	15.4	27.1	44.7
FSVD	<i>M</i> = 4	0.0	12.4	42.0	71.9	100.0	0.0	4.4	15.4	27.0	44.5
	<i>M</i> = 6	0.0	13.3	43.1	73.0	100.0	0.0	4.5	15.6	27.2	44.6
	<i>M</i> = 8	0.0	13.3	43.2	73.1	100.0	0.0	4.3	15.4	27.0	44.7
<i>T = 20</i>											
DSS-Q		0.0	23.8	56.3	88.7	100.0	0.0	11.2	25.3	40.2	62.5
PSVD-Q	<i>M</i> = 4	0.0	23.9	56.3	88.5	100.0	0.0	11.6	25.7	40.0	59.9
	<i>M</i> = 6	0.0	24.2	56.4	88.4	100.0	0.0	11.2	25.3	40.2	62.7
	<i>M</i> = 8	0.0	23.8	56.2	88.6	100.0	0.0	11.2	25.5	40.4	62.6
FSVD	<i>M</i> = 4	0.0	23.4	56.5	89.2	100.0	0.0	10.6	25.1	40.5	63.3
	<i>M</i> = 6	0.0	24.3	57.2	89.7	100.0	0.0	10.5	24.9	40.4	63.4
	<i>M</i> = 8	0.0	24.4	57.2	89.7	100.0	0.0	10.7	25.1	40.4	63.2
<i>T = 30</i>											
DSS-Q		0.0	32.0	66.9	100.0	100.0	0.0	18.0	35.4	53.7	79.8
PSVD-Q	<i>M</i> = 4	0.0	31.8	66.8	100.0	100.0	0.0	17.8	35.0	51.5	70.3
	<i>M</i> = 6	0.0	31.8	66.9	100.0	100.0	0.0	17.8	35.3	53.7	77.6
	<i>M</i> = 8	0.0	31.9	67.1	100.0	100.0	0.0	17.4	35.6	54.7	80.0
FSVD	<i>M</i> = 4	0.0	32.3	67.8	100.0	100.0	0.0	17.9	35.3	53.8	80.3
	<i>M</i> = 6	0.0	32.9	68.3	100.0	100.0	0.0	17.4	35.3	54.1	80.7
	<i>M</i> = 8	0.0	32.8	68.4	100.0	100.0	0.0	17.5	35.2	54.0	80.7
<i>T = 40</i>											
DSS-Q		0.0	38.0	76.8	100.0	100.0	0.0	24.3	44.5	65.4	94.1
PSVD-Q	<i>M</i> = 4	0.0	37.9	76.8	100.0	100.0	0.0	23.9	43.0	59.9	74.9
	<i>M</i> = 6	0.0	37.9	76.8	100.0	100.0	0.0	24.3	44.4	64.8	86.6
	<i>M</i> = 8	0.0	37.8	76.9	100.0	100.0	0.0	24.3	44.6	65.5	92.3
FSVD	<i>M</i> = 4	0.0	38.3	77.2	100.0	100.0	0.0	24.3	44.6	66.1	94.7
	<i>M</i> = 6	0.0	38.7	77.6	100.0	100.0	0.0	24.2	44.5	65.9	94.6
	<i>M</i> = 8	0.0	38.8	77.6	100.0	100.0	0.0	24.1	44.5	65.7	94.6

The table reports the optimal allocation to the risky asset obtained by the quadrature-based discretized state space method (referred to as DSS-Q), the partial SVD approach that uses quadrature (referred to as PSVD-Q), and the full SVD approach (referred to as FSVD). The preferences are CRRA and the returns predictable by the log dividend yield. The data generating process and the parameters used are described in Eqs. 13 and 27, respectively. The annualized risk-free rate is set equal to 6%. We consider two values for the coefficient of relative risk aversion, $\gamma = 5$, and 15 and four values for the time horizon, $T = 10, 20, 30$, and 40. We present the results for 5 different initial values of the log dividend yield z . Each value corresponds to the p -th percentile of the unconditional distribution of z , denoted by z_p , where p takes values 10, 30, 50, 70, and 90.

equivalent return does not exceed a few basis points and that the absolute difference in optimal allocations is not more than 2% compared to the quadrature method, when M is sufficiently large, i.e., $M = 6$ or 8. It is important to emphasize that the SVD approach performs extremely well over the entire state space of the predictor, a property that is crucial for addressing realistic problems with long horizons and multiple

assets/state variables. It is worth emphasizing that, in contrast, vBB report results *only* for the mean of the predictor. In summary, the evidence we produce in this subsection provides strong support for the overall accuracy of the SVD approach that uses iteration on the certainty equivalent function and Taylor approximations.

5 Conclusion

In this paper, we address the issue of accuracy of the value function iteration (VFI) approach in the context of numerical solutions to portfolio choice problems. A recent paper by [van Binsbergen and Brandt \(2007\)](#) argues that the VFI approach is inferior to the portfolio weight iteration approach developed by [Brandt et al. \(2005\)](#), especially for high levels of risk aversion and long investment horizon. We argue that this conclusion applies only to the specific way that VFI is implemented by [van Binsbergen and Brandt \(2007\)](#) and that the problem arises because [van Binsbergen and Brandt \(2007\)](#) approximate the highly nonlinear value function by low-order polynomials. We demonstrate that working with the certainty equivalent version of the value function can address this problem. To illustrate our argument, we show that, for a portfolio choice problem with predictable returns and CRRA preferences, the state variable decomposition approach, which uses value function iteration on the certainty equivalent function as recently proposed by [Garlappi and Skoulakis \(2008\)](#), produces very accurate results over the entire state space.

Acknowledgements We are grateful to Pete Kyle, Mark Loewenstein, and Jacob Sagi for helpful discussions and comments.

References

- Balduzzi, P., & Lynch, A. (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.
- Brandt, M. W. (2004). Portfolio choice problems. In Y. Ait-Sahalia, & L. P. Hansen (Eds.), *Handbook of financial econometrics*, Elsevier, 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.
- Brennan, M. J., Schwartz, E. S., & Lagnado, R. (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control*, *21*, 1377–1403.
- Campbell, J. Y., & Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics*, *114*, 433–495.
- Cox, J. C., & Huang, C. (1989). Optimal portfolio and consumption policies when asset prices follow a diffusion process. *Journal of Economic Theory*, *49*, 33–83.
- Detemple, J. B., Garcia, R., & Rindisbacher, M. (2003). A Monte Carlo method for optimal portfolios. *Journal of Finance*, *58*, 410–446.
- Garlappi, L., & Skoulakis, G. (2008). A state variable decomposition approach for solving portfolio choice problems. Working paper, University of Texas at Austin.
- Judd, K. (1998). *Numerical methods in economics*, Cambridge: MIT Press.
- Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, *14*, 113–147.
- Merton, R. C. (1973). An intertemporal asset pricing model. *Econometrica*, *41*, 867–888.

- Savits, T. H. (2006). Some statistical applications of Faà di Bruno. *Journal of Multivariate Analysis*, 97, 2131–2140.
- Tsitsiklis, J. N., & Van Roy, B. (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12, 694–703.
- 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.
- Willink, R. (2005). Normal moments and Hermite polynomials. *Statistics and Probability Letters*, 73, 271–275.