

A State Variable Decomposition Methodology for Solving Portfolio Choice Problems*

Lorenzo Garlappi[†]

Georgios Skoulakis[‡]

First draft: November 2007

This draft: December 2008

*We are grateful to Gurdip Bakshi, Ravi Jagannathan, Raymond Kan, Pete Kyle, Mark Loewenstein, Oleg Rytchkov, Jacob Sagi, Costis Skiadas, Stathis Tompaidis, Raman Uppal and seminar participants at the 2007 International Conference on Computing in Economics and Finance, Montréal, the 2008 Western Finance Association meetings, Hawaii, the 2008 European Finance Association meetings, Athens, and the University of Texas at Austin for helpful discussions and comments. Any errors are our sole responsibility. Comments and suggestions are welcome at lorenzo.garlappi@mcombs.utexas.edu or gskoulak@rhsmith.umd.edu.

[†]Finance Department, McCombs School of Business, University of Texas at Austin, Austin TX, 78712; Email: lorenzo.garlappi@mcombs.utexas.edu.

[‡]Finance Department, Robert H. Smith School of Business, University of Maryland, College Park, MD 20742; Email: gskoulak@rhsmith.umd.edu.

A State Variable Decomposition Methodology for Solving Portfolio Choice Problems

Abstract

This paper develops a new method for the solution of a broad class of discrete-time dynamic portfolio choice problems. The method efficiently approximates conditional expectations of the value function by using (i) a *decomposition* of the state variables into a predictable component and a stochastic deviation, and (ii) a *Taylor expansion* of the value function. The outcome of this *State Variable Decomposition (SVD)* approach is an approximate problem in which conditional expectations can be computed efficiently and, under standard distributional assumptions, even *analytically*, without sacrificing precision. We illustrate the accuracy of the SVD method in handling several realistic features of portfolio choice problems such as intermediate consumption, multiple risky assets, multiple state variables, portfolio constraints, non-time-separable preferences, and non-redundant endogenous state variables. We finally use the SVD method to solve a large-scale life-cycle portfolio problem with predictable returns, recursive preferences and realistic portfolio constraints. The versatility of our approach makes it a suitable toolkit for a wide range of dynamic problems in finance and economics.

Keywords: Portfolio choice, numerical methods, Taylor series approximations

JEL Classification: G11, G12

1 Introduction

The formation and management of portfolios of financial assets for long-lived investors are among the most fundamental economic decisions faced by both individual and institutions. The study of such issues has a long and celebrated history in financial economics, starting with the seminal work of Samuelson (1969), and Merton (1969, 1971). With a few notable exceptions,¹ the vast majority of realistic dynamic portfolio choice problems are analytically intractable and their solution usually relies on complex numerical optimization procedures. The computationally intensive nature of existing numerical methods severely restricts the range of potential applications to problems with a small number of assets and/or state variables. Moreover, numerical complexity limits the empirical relevance of optimal portfolio choice models as it renders any serious calibration to real data practically infeasible.

The objective of this paper is to develop a simple, precise, and efficient numerical methodology for the solution of a large class of discrete-time dynamic portfolio choice problems, and, more generally, for a host of dynamic problems in economics and finance that are characterized by a similar recursive structure. The method we propose relies on accurately approximating conditional expectations, which are the quantities that lie at the heart of every stochastic dynamic programming problem. Our methodology delivers an approximation that is easily computed and, under standard distributional assumptions, available in *closed form*. The method does not require the use of computationally intensive numerical integration techniques such as quadrature or Monte-Carlo methods and therefore facilitates computational efficiency. More importantly, it does so without sacrificing precision, as we illustrate through a number of examples.

Our approximation scheme rests on two simple building blocks. The first is a *decomposition* of each state variable characterizing the problem into a predictable component and the associated stochastic deviation. The second is the use of a *Taylor expansion* for approximating the value function of the problem by a polynomial in the stochastic deviations. These two simple steps allow us to approximate the next-period value function by a sum of products of two separate functions: one depending only on stochastic shocks and one depending only on choice variables. This separation between shocks and choice variables is computationally efficient because it reduces the complex problem of computing conditional expectations of the value function to the much simpler problem of computing the moments of the shocks driving the state variables. Since the above decomposition of the state variables is key for achieving this separation, we use the term *State Variable Decomposition* (SVD hereafter) to refer to our approach.

¹Obtaining a fully closed-form solution to a portfolio choice/consumption problem typically requires particular assumptions about preferences, market completeness and absence of frictions and constraints. Examples of recent papers that derive closed-form solutions include Kim and Omberg (1996), Wachter (2002), Chacko and Viceira (2005), Cvitanić, Lazrak, Martellini, and Zapatero (2006), Liu (2007), and Aït-Sahalia, Cacho-Diaz, and Hurd (2008).

Before applying any approximation scheme, it is however important to be aware of the possible sources of the error incurred. There are two main sources of error that can arise from applying the SVD method. The first is the *projection error* caused by the fact that the value function of the problem is usually not known analytically and therefore has to be approximated over the state space. In practice, the value function used in the numerical solution of a dynamic programming problem is in fact a *projection* of the true unknown value function on a space generated by a set of basis functions (e.g., polynomials). The projection error is common to various numerical methods used to solve dynamic programming problems. The second source of error, unique to the SVD method, is the *Taylor error* which is incurred by using Taylor expansions to approximate the (projection of the) value function as a polynomial in the stochastic deviations to the state variables.

In implementing the SVD approach, we take measures to minimize both sources of approximation error. To minimize the projection error, instead of working with the value function directly, we work with the *certainty equivalent* function which is a strictly monotonic transformation of the value function containing the same information. As Garlappi and Skoulakis (2008a) document, the certainty equivalent is a much less nonlinear function than the value function, and allows for a more efficient and accurate approximation through projection methods. To minimize the Taylor error, we carefully select the type of the Taylor expansion we use to guarantee the convergence of the method as the number of Taylor terms increases. Garlappi and Skoulakis (2008b) formally address the issue of convergence of Taylor approximations in the context of a static portfolio choice problem with HARA preferences. We rely on the theoretical results of this paper in the implementation of the SVD method.

As stated above, the essence of the SVD method is to reduce the original problem into an approximate one that involves conditional expectations that are functions *only* of the stochastic deviations to the state variables, i.e., conditional expectations do not depend on choice variables. In several cases, e.g., when asset returns and state variables follow normal or lognormal distributions, such expectations can be obtained analytically in closed form. Alternatively, it is always possible to use either quadrature methods (as in Judd (1998)) or simulation/regression methods (as in Longstaff and Schwartz (2001), Tsitsiklis and Van Roy (2001), and Brandt, Goyal, Santa-Clara, and Stroud (2005)). The crucial advantage of our decomposition is that, regardless of the method used to compute conditional expectations in the *approximate* problem, this computation needs not be repeated each time a candidate solution is considered in the numerical optimization process. This repetition is responsible for the loss of efficiency of traditional quadrature methods applied to the *original* problem, especially when the number of state variables is large. In contrast, our method is capable of efficiently handling problems with a large number of state variables, general return

distributions, constraints and, most importantly, problems with non-redundant *endogenous* state variables (e.g., wealth).

There is an extensive literature on numerical and analytical approximations to the solution of portfolio choice problems, starting from the early work on polynomial approximations to the investor’s utility (see, for example, Samuelson (1970), Hakansson (1971), Loistl (1976), Pulley (1981, 1983), Kroll, Levy, and Markowitz (1984), Markowitz (1991), and Hlawitschka (1994)) and continuing, more recently, with methods based on numerical solution of PDEs (e.g., Brennan, Schwartz, and Lagnado (1997)); log-linearization of the budget constraint (Campbell and Viceira (1999), Campbell, Chan, and Viceira (2003)); perturbation and projection methods (Kogan and Uppal (2003) and Das and Sundaram (2002)); state space discretization (Barberis (2000), Balduzzi and Lynch (1999), Brandt (1999)); and, Malliavin calculus and Monte Carlo methods (Detemple, Garcia, and Rindisbacher (2003)). We contribute to the literature by proposing a method that can handle a broad array of challenging and economically relevant features. In particular, our method allows for a large number of assets and state variables, portfolio constraints, time-separable (including non-homothetic) as well as recursive preferences, and does not require market completeness.

A paper closely related to ours is Brandt, Goyal, Santa-Clara, and Stroud (2005) (BGSS hereafter) who rely on simulation and regression methods for the computation of conditional expectations of the value function. Like SVD, the BGSS method uses Taylor approximations of the value function to simplify the computation of these expectations. Our approach, however, differs from BGSS in several important aspects. First, the SVD method is based on *value function iteration* while BGSS is based on *policy function iteration*. As we show in Section 4.2, a method that relies on policy function iteration in conjunction with Taylor approximation cannot handle problems in which the choice variable depends in a non-trivial fashion on endogenous state variables such as wealth. This is not the case for the SVD method which can easily handle endogenous state variables. Second, under standard distributional assumptions for the fundamental shocks (e.g., normal and log-normal distributions) the SVD method allows the computation of conditional expectation analytically, without the need to rely on simulation and regression methods. Third, the SVD method allows total flexibility in the construction of the Taylor expansion used to approximate the value function. This aspect of the method is crucial for guaranteeing the convergence of the Taylor series approximation, as the number of Taylor terms increases. In our applications, we document that the choice of the center of expansion can have a substantial effect on computational accuracy, especially when solving problems with multiple risky assets.

In summary, the main advantages of the method we propose are: (i) conceptual simplicity: we rely on a state variable decomposition and simple Taylor series approximations; (ii) computational convenience:

the particular approximation used allows efficient computation of the required conditional expectations, either analytically or numerically; (iii) generality: the method is applicable to a variety of problems that involve dynamic optimization, including problems in which the optimal portfolio is dependent on wealth; (iv) precision: the flexibility in the construction of the Taylor series helps to minimize and control the approximation error. Although the approach we propose is general enough to be applied to a large class of dynamic programming problems, in this paper we focus explicitly on portfolio choice problems to facilitate the comparison of our method to existing ones that have been proposed in the finance literature.

To assess the performance of the SVD method, we apply it to several static and dynamic portfolio choice problems. We compare the SVD solution to the exact analytical solution whenever the latter is available. Alternatively, we use as a benchmark for comparison the approximate numerical solution obtained via traditional quadrature-based methods. It is important to note, however, that quadrature-based solutions are largely computationally inefficient and practically feasible only for small-scale problems.

In all the applications we consider, the SVD method proves to be extremely accurate, producing solutions that are virtually indistinguishable from the benchmark used for comparison. From these applications, we deduce two important methodological points that are relevant for understanding the potential error of a numerical approximating scheme that relies on Taylor expansions like ours. First, we show that the choice of the center of expansion is extremely important especially when solving problems with multiple assets. In particular, an ill-chosen center of expansion can lead to severe losses in terms of certainty equivalent return. The SVD method, by appropriately choosing the center of expansion, does not suffer from this shortcoming and guarantees convergence and minimal losses in accuracy.

The second methodological point that emerges from our applications is specific to the case of dynamic portfolio choice problems. Our analysis shows that the use of Taylor expansions in solving approximate portfolio choice problems in which wealth is not a redundant state variable cannot abstract away from the knowledge of the value function and its properly computed derivatives. While the SVD approach is accurate and always converges as the number of terms in the approximation increases, methods that rely on policy iteration and Taylor expansions, like BGSS, are bound to ignore the dependence of the portfolio weights on wealth. The consequences of ignoring this dependence can be detrimental to the accuracy of the approximation. In fact, in the context of a simple multiperiod problem with CARA utility and normally distributed asset returns, we show that overlooking the dependence of portfolio weights on wealth may lead to *divergence* of the Taylor approximation.

We conclude our analysis of the SVD methodology by applying it to a realistic life-time multi-asset portfolio choice problem of a finitely-lived investor with recursive preferences who faces a time-varying

investment opportunity set. Following Campbell, Chan, and Viceira (2003), we consider a menu of three risky assets (nominal T-bill, nominal T-bond, and equity) whose risk premia are assumed to be predictable by the dividend-yield, the term spread, the short nominal interest as well as by lagged values of the asset returns themselves. Our approach allows us to impose realistic borrowing and short-selling constraints yielding a challenging dynamic problem with *six* state variables and *three* choice variables. Despite the complexity of this problem, the SVD approach produces results that are indistinguishable from those obtained from more traditional, but also much less efficient, quadrature-based methods.² More importantly, the SVD method can be employed to solve problems with recursive utility even when the elasticity of intertemporal substitution (EIS) is not close to 1, *without* any additional computational effort. In contrast, analytical approximation methods such as the log-linearization approach of Campbell and Viceira (1999), delivers accurate results only for short time intervals and values of EIS close to unity.

The rest of the paper proceeds as follows. In Section 2, we develop the SVD approach. Section 3 illustrates the precision of the approach by applying it to standard static portfolio choice problems with CRRA preferences. In Section 4, we apply the SVD method to dynamic portfolio choice problems with CARA preferences and analyze its convergence properties in comparison to the BGSS method. In Section 5, we apply the SVD method to solve a strategic asset allocation problem with predictable returns. Section 6 concludes. Appendix A describes efficient recursive schemes for the computation of moments of the multivariate normal distribution and derivatives of composite functions of multiple variables. Appendix B contains proofs of propositions while in Appendix C we present tables with numerical results.

2 The State Variable Decomposition (SVD) approach

The main idea behind the SVD approach to solving portfolio choice problems is to approximate the original problem by one in which portfolio weights are *separated*—in a sense to be made precise later—from the shocks to the state variables describing the evolution of the problem over time. This separation is achieved by first decomposing each state variable into a predictable component and a stochastic deviation, and then approximating the original problem via a Taylor expansion with respect to the stochastic deviation. A direct consequence of the separation between portfolio weights and shocks is that the conditional expectations that need to be computed for solving the portfolio problem *do not depend* on portfolio weights. As we document extensively in Sections 3–5, this feature significantly enhances the computational efficiency without sacrificing the accuracy of the solution.

²In the solution of a life-cycle portfolio choice and consumption problem with recursive preferences, the SVD method was found to be 35 times faster than the quadrature method. See Section 5 for further details.

In the next subsection we illustrate the SVD method in a simple problem of an investor with CRRA preferences who maximizes expected utility of terminal wealth when the investment opportunity set is stochastic. In Subsection 2.2, we present the general version of the SVD method.

2.1 An illustrative example

Consider an investor with investment horizon T and terminal utility $U(\cdot)$ of the CRRA form, $U(W_T) = W_T^{1-\gamma}/(1-\gamma)$ where W_T denotes the terminal wealth, and $\gamma > 1$ is the coefficient of relative risk aversion. The investment opportunity set consists of *two* risky assets and a risk-free asset.³ The investor's problem can be described as

$$\max_{\{\boldsymbol{\omega}_t\}_{t=0}^{T-1}} E_0[U(W_T)], \quad (1)$$

where $\boldsymbol{\omega}_t = (\omega_{1,t}, \omega_{2,t})'$ is the two-dimensional portfolio weight vector on the risky assets at time t . The wealth evolves according to the budget equation

$$W_{t+1} = W_t(R_f + \boldsymbol{\omega}'_t \mathbf{R}_{t+1}), \quad (2)$$

where \mathbf{R}_{t+1} is the two-dimensional vector of asset returns in excess of the (gross) risk-free rate R_f . We assume that the investment opportunity set is stochastic and characterized by an *exogenous* state variable \mathbf{S}_t . A state variable is said to be exogenous if it does not depend on the choice variable $\boldsymbol{\omega}_t$. By contrast, the wealth W_t is an *endogenous* state variable since it does depend on the choice variable $\boldsymbol{\omega}_t$. The information available to the investor at time t is summarized by the state variables W_t and \mathbf{S}_t .

The investor's problem is characterized by the sequence of problems

$$J_t(W_t, \mathbf{S}_t) = \max_{\{\boldsymbol{\omega}_s\}_{s=t}^{T-1}} E_t[U(W_T)], \quad t = 0, \dots, T-1 \quad (3)$$

subject to the budget constraint (2), where $J_t(\cdot, \cdot)$ is the so-called *value function* at time t . By the Bellman principle of optimality, (3) is equivalent to the following recursion

$$J_t(W_t, \mathbf{S}_t) = \max_{\boldsymbol{\omega}_t} E_t [J_{t+1}(W_t(R_f + \boldsymbol{\omega}'_t \mathbf{R}_{t+1}), \mathbf{S}_{t+1})], \quad (4)$$

where the terminal condition is given by $J_T(W_T, \mathbf{S}_T) = U(W_T) = \frac{W_T^{1-\gamma}}{1-\gamma}$. The solution of a dynamic programming problem traditionally involves backward iteration on the value function J_t . Typically, in each step of the backward recursion, the value function cannot be characterized analytically and has to be

³Without loss of generality, we restrict the discussion to two risky assets to reduce notational complexity and keep the exposition clear. The analysis extends to the case of multiple risky assets in a straightforward fashion.

approximated via, for example, state space discretization with interpolation or projection methods.⁴ To this end, one constructs a grid in the state space and computes the value function at the grid points. This information is then used to provide an approximation to the value function over the entire state space. Due to the dynamic nature of the problem, it is important to ensure that the approximation to the value function is sufficiently accurate in each step of the backward recursion.

As a preparatory step for the application of the SVD method, we note that one can perform the backward recursion based on a strictly monotonic transformation of $J_t(\cdot)$, instead of $J_t(\cdot)$ itself, with no loss of information. Clearly, working with such a transformation leaves the solution of the problem unaffected. Motivated by the economic structure of the problem under consideration, a natural choice for this transformation is the inverse of the utility function $U^{-1}(\cdot)$. In this case, the transformed value function $V_t = U^{-1}(J_t)$ is simply the *certainty equivalent* of J_t , defined implicitly by

$$J_t(W_t, \mathbf{S}_t) = U(V_t(W_t, \mathbf{S}_t)). \quad (5)$$

As Garlappi and Skoulakis (2008a) show, in the context of a portfolio choice problem with CRRA preferences and predictable returns, the certainty equivalent function $V_t(\cdot)$ turns out to be much less nonlinear than the value function $J_t(\cdot)$ and, therefore, much easier to approximate. The advantage of using the certainty equivalent function is more pronounced when the utility function is highly concave, that is for high levels of risk aversion.

Given the homotheticity property of CRRA preferences, it is well-known that the value function $J_t(\cdot, \cdot)$ is multiplicatively *separable* in its two arguments, W_t and \mathbf{S}_t , i.e., there exists a function $\mathcal{V}_t(\mathbf{S}_t)$ such that the original value function can be written as follows

$$J_t(W_t, \mathbf{S}_t) = \frac{W_t^{1-\gamma}}{1-\gamma} \mathcal{V}_t(\mathbf{S}_t)^{1-\gamma}. \quad (6)$$

The above decomposition implies that the certainty equivalent function defined in (5) is separable in its two arguments:

$$V_t(W_t, \mathbf{S}_t) = W_t \mathcal{V}_t(\mathbf{S}_t). \quad (7)$$

Due to the structure of equation (7), we refer to the function $\mathcal{V}_t(\cdot)$ as the *certainty equivalent return* function. Note that the function $\mathcal{V}_t(\cdot)$ depends *only* on the exogenous state variables \mathbf{S}_t . Substituting (6) in (4) and using the definition of certainty equivalent in (5), we arrive at the following Bellman equation

$$U(W_t \mathcal{V}_t(\mathbf{S}_t)) = \max_{\omega_t} E_t [U(W_{t+1} \mathcal{V}_{t+1}(\mathbf{S}_{t+1}))], \quad (8)$$

⁴See Judd (1998) for a textbook treatment of value function approximation in the context of dynamic programming problems.

with terminal condition $\mathcal{V}_T(\mathbf{S}_T) = 1$. This is the starting point for the application of the SVD methodology.

As stated earlier, the key idea of the SVD method is to *separate* the choice variable $\boldsymbol{\omega}_t$ from the shocks to the state variable. Specifically, this amounts to approximating the argument in the conditional expectation on the right-hand side of (8) by a *sum of products* of quantities that depend only on the choice variables $\boldsymbol{\omega}_t$, and quantities that depend only on the fundamental shocks to the state variables. To achieve this separation, we first decompose the state variables into a predictable component and a stochastic deviation and then apply a Taylor expansion with respect to this deviation. As a consequence, the conditional expectations in the resulting approximate problem will depend only on shocks to the state variables and not on choice variables. This separation considerably simplifies the computation of the expectations that can then be easily carried out either analytically or numerically.

Once this separation is achieved, the method relies on backward recursion in the standard fashion, starting from the solution of the problem at the final date T and proceeding until time zero. In the illustration that follows, we detail the steps required by the SVD method at a generic time t . To this purpose, we assume that an approximation to the certainty equivalent return function $\mathcal{V}_{t+1}(\cdot)$ has already been obtained and show how to obtain the optimal allocation $\boldsymbol{\omega}_t$ and the function $\mathcal{V}_t(\cdot)$ at time t .

We now describe in detail the three steps required for the implementation of the SVD method. While, in principle, the SVD approach is based on a decomposition of *all* state variables, an operational variant of the method decomposes only the endogenous state variable, e.g., wealth. In this subsection, we focus on this variant for purposes of illustration. The general SVD method is described in the next subsection.

Step 1. Decomposition of wealth

In the first step of the SVD approach, we decompose the risky asset return \mathbf{R}_{t+1} into a predictable component $\mathbf{c}_{R,t} = (c_{1,t}, c_{2,t})'$, known at time t , and the associated stochastic deviation $\boldsymbol{\varepsilon}_{R,t+1} = (\varepsilon_{R_1,t+1}, \varepsilon_{R_2,t+1})'$: $\mathbf{R}_{t+1} = \mathbf{c}_{R,t} + \boldsymbol{\varepsilon}_{R,t+1}$. This return decomposition gives rise to the following decomposition of wealth

$$W_{t+1} = W_t(R_f + \boldsymbol{\omega}'_t \mathbf{R}_{t+1}) = c_{W,t}(\boldsymbol{\omega}_t) + \varepsilon_{W,t+1}(\boldsymbol{\omega}_t), \quad (9)$$

where $c_{W,t}(\boldsymbol{\omega}_t) = W_t(R_f + \boldsymbol{\omega}'_t \mathbf{c}_{R,t})$ is the predictable wealth component and $\varepsilon_{W,t+1}(\boldsymbol{\omega}_t) = W_t(\boldsymbol{\omega}'_t \boldsymbol{\varepsilon}_{R,t+1})$ is the corresponding stochastic deviation.

Step 2. Separation of choice variables from shocks

The second step involves separating the choice variable $\boldsymbol{\omega}_t$ from the shocks $\boldsymbol{\varepsilon}_{R,t+1}$ in order to express the argument in the conditional expectation of the Bellman equation (8) as a sum of products of quantities

that depend only on the choice variable $\boldsymbol{\omega}_t$ and quantities that depend only on the shocks $\boldsymbol{\varepsilon}_{R,t+1}$. Using the functional form of $U(\cdot)$ and the decomposition of wealth in (9), the conditional expectation in (8) becomes

$$E_t [U(W_{t+1}\mathcal{V}_{t+1}(\mathbf{S}_{t+1}))] = E_t [U(c_{W,t}(\boldsymbol{\omega}_t) + \varepsilon_{W,t+1}(\boldsymbol{\omega}_t))\mathcal{V}_{t+1}(\mathbf{S}_{t+1})^{1-\gamma}]. \quad (10)$$

Using the M -th order Taylor approximation of the quantity $U(c_{W,t}(\boldsymbol{\omega}_t) + \varepsilon_{W,t+1}(\boldsymbol{\omega}_t))$ with respect to the stochastic deviation $\varepsilon_{W,t+1}(\boldsymbol{\omega}_t)$ around 0, we obtain

$$U(c_{W,t}(\boldsymbol{\omega}_t) + \varepsilon_{W,t+1}(\boldsymbol{\omega}_t)) \approx \sum_{m=0}^M \frac{1}{m!} U^{(m)}(c_{W,t}(\boldsymbol{\omega}_t)) [\varepsilon_{W,t+1}(\boldsymbol{\omega}_t)]^m, \quad (11)$$

where $U^{(m)}(c_{W,t}(\boldsymbol{\omega}_t))$ denotes the m -th derivative of $U(\cdot)$ evaluated at $c_{W,t}(\boldsymbol{\omega}_t)$.

The Taylor approximation in this step is similar to the one used in BGSS. The only crucial difference is that the BGSS method uses $c_{W,t} = W_t R_f$ as the center of the Taylor expansion, while the SVD method allows for arbitrary choice of $c_{W,t}$. It turns out that the choice of the center of expansion is critical for ensuring convergence of the Taylor series approximation, as we document in Section 3. Loistl (1976) and Hlawitschka (1994) study the theoretical and empirical convergence properties of Taylor series approximation to expected utility and raise concerns about the validity of the approach. Garlappi and Skoulakis (2008b) offer a thorough theoretical analysis of the subject and provide conditions under which the Taylor series approximation does converge in the context of expected utility computations as well as static optimal portfolio choice problems.

To achieve full separation between the choice variables $\boldsymbol{\omega}_t$ and the stochastic deviation $\boldsymbol{\varepsilon}_{R,t+1}$ in (11), we use the binomial formula to express the quantity $[\varepsilon_{W,t+1}(\boldsymbol{\omega}_t)]^m = W_t^m (\boldsymbol{\omega}'_t \boldsymbol{\varepsilon}_{R,t+1})^m$ as follows

$$[\varepsilon_{W,t+1}(\boldsymbol{\omega}_t)]^m = W_t^m \sum_{k=0}^m \frac{m!}{k!(m-k)!} \left(\omega_{1,t}^k \omega_{2,t}^{m-k} \right) \cdot \left(\varepsilon_{R_1,t+1}^k \varepsilon_{R_2,t+1}^{m-k} \right). \quad (12)$$

It follows from the functional form of U that the derivatives $U^{(m)}(c_{W,t}(\boldsymbol{\omega}_t))$ in (11) are

$$U^{(m)}(c_{W,t}(\boldsymbol{\omega}_t)) = W_t^{1-\gamma-m} U^{(m)}(c_{p,t}(\boldsymbol{\omega}_t)), \quad (13)$$

where $c_{p,t}(\boldsymbol{\omega}_t) = R_f + \boldsymbol{\omega}'_t \mathbf{c}_{R,t}$ is the predictable portfolio return component. Combining equations (11), (12), and (13) yields

$$U(c_{W,t}(\boldsymbol{\omega}_t) + \varepsilon_{W,t+1}(\boldsymbol{\omega}_t)) \approx W_t^{1-\gamma} \sum_{m=0}^M \sum_{k=0}^m \frac{1}{k!(m-k)!} U^{(m)}(c_{p,t}(\boldsymbol{\omega}_t)) \left(\omega_{1,t}^k \omega_{2,t}^{m-k} \right) \cdot \left(\varepsilon_{R_1,t+1}^k \varepsilon_{R_2,t+1}^{m-k} \right). \quad (14)$$

Note that equation (14) is of the desired form: the right-hand side is expressed as a sum of products of quantities that depend only on the choice variable $\boldsymbol{\omega}_t$ and quantities that depend only on the stochastic

deviations $\varepsilon_{R,t+1}$. Using equations (10) and (14) and eliminating the term $W_t^{1-\gamma}$ from both sides of the original Bellman equation (8), we obtain the following *approximate* Bellman equation⁵

$$U(\mathcal{V}_t(\mathbf{S}_t)) = \max_{\boldsymbol{\omega}_t} \sum_{m=0}^M \sum_{k=0}^m F_{t,m,k}(\boldsymbol{\omega}_t) E_t [G_{t+1,m,k}(\varepsilon_{R,t+1}) \mathcal{V}_{t+1}(\mathbf{S}_{t+1})^{1-\gamma}], \quad (15)$$

where

$$F_{t,m,k}(\boldsymbol{\omega}_t) = \frac{1}{k!(m-k)!} U^{(m)}(c_{p,t}(\boldsymbol{\omega}_t)) \left(\omega_{1,t}^k \omega_{2,t}^{m-k} \right), \quad (16)$$

$$G_{t+1,m,k}(\varepsilon_{R,t+1}) = \varepsilon_{R_1,t+1}^k \varepsilon_{R_2,t+1}^{m-k}. \quad (17)$$

For any given portfolio $\boldsymbol{\omega}_t$, the terms $F_{t,m,k}(\boldsymbol{\omega}_t)$ can be readily obtained using the predictable component of risky asset returns, i.e., $\mathbf{c}_{R,t}$.

To solve for the optimal portfolio $\boldsymbol{\omega}_t$, it is necessary to compute the conditional expectations in (15). As a consequence of the previous two steps, the functions inside these expectations do not depend on the choice variables $\boldsymbol{\omega}_t$ but only on the return shocks $\varepsilon_{R,t+1}$ and the exogenous state variable \mathbf{S}_{t+1} . From a computational point of view, this feature is extremely convenient. In fact, in the optimization process used to find $\boldsymbol{\omega}_t$ the conditional expectations in (15) need to be computed only *once* for each grid point in the state space leading to significant reduction in computation.

Step 3. Computation of conditional expectations

The third, and final, step of our methodology concerns the computation of the conditional expectations in the approximate Bellman equation (15). In the context of the illustrative example of this subsection, this step can be accomplished numerically following two different approaches.⁶

The first approach is to use quadrature methods as in Judd (1998) to compute the conditional expectations in (15). This is particularly convenient especially when the shocks to the returns $\varepsilon_{R,t+1}$ and the exogenous state variables \mathbf{S}_{t+1} follow standard distributions such as normal. As discussed above, the computation of these expectations is done only once for each grid point in the state space. We refer to this approach as the SVD method with expectations computed via *quadrature* (SVD-Q hereafter). Note that this method differs from the traditional quadrature-based approach to the solution of portfolio choice problems where the conditional expectations in the Bellman equation (8) need to be computed each time a new candidate portfolio $\boldsymbol{\omega}_t$ is evaluated in the optimization process. It is well-known that the traditional

⁵Note the slight abuse of notation, because we use the same symbol \mathcal{V}_t for both the exact and the approximate certainty equivalent return function.

⁶In the next subsection, where we develop a general version of the SVD method, we show how, under standard distributional assumptions, these conditional expectations can be computed analytically.

quadrature-based approach suffers from a severe curse of dimensionality. In contrast, the SVD-Q method requires computation of conditional expectations only once for each grid point and, therefore, is much more efficient than the traditional quadrature-based approach.

The second approach to computing the conditional expectations in (15) is to rely on the simulation-based methodology of parameterized expectations of Tsitsiklis and Van Roy (2001) and Longstaff and Schwartz (2001). This is also the method followed by BGSS in computing expectations. We refer to this approach as the SVD method with expectations computed via *simulation* (SVD-S hereafter).⁷

Once the conditional expectations are computed using one of the previous two methods, we can determine the optimal allocation ω_t and, consequently, the certainty equivalent return function $\mathcal{V}_t(\cdot)$ for each grid point in the state space. The procedure is then repeated until we reach time zero.

In the preceding analysis, wealth is the only state variable to be decomposed and used in the Taylor expansion. While this captures the essence of the SVD approach, one, alternatively, could also decompose the additional exogenous state variables \mathbf{S}_{t+1} and employ a Taylor approximation with respect to *all* state variables. This variant of the SVD approach turns out to be suitable not only for the homothetic type of portfolio choice problem considered in this subsection but also for more general problems. In the remainder of this section, we relax the homotheticity assumption of the CRRA case and show how the SVD method can be applied to handle a wide class of portfolio choice problems.

2.2 The general case

The general form of the dynamic portfolio choice problems we study can be described as follows:

$$J_t(\mathbf{S}_t) = \max_{\omega_t} \{ \mathcal{H}_t(u_t(\mathbf{S}_t, \omega_t), E_t[J_{t+1}(\mathbf{S}_{t+1})]) \}, \text{ where } \mathbf{S}_{t+1} = \Gamma_t(\mathbf{S}_t, \omega_t, \delta_{t+1}). \quad (18)$$

In the above expressions, \mathbf{S}_t denotes a d_S -dimensional vector of state variables whose evolution through time is described by the law of motion $\Gamma_t(\cdot)$; ω_t is a d_ω -dimensional vector of choice variables; and δ_{t+1} is a d_δ -dimensional vector of innovations. We allow the innovation term to be a function of the state variable \mathbf{S}_t and a vector of fundamental shocks ε_{t+1} , that is, $\delta_{t+1} = \delta_{t+1}(\mathbf{S}_t, \varepsilon_{t+1})$.⁸ The state variable vector \mathbf{S}_t is composed of both endogenous and exogenous state variables. By definition, the value of an endogenous

⁷As we discuss more in detail in Section 4.2, our method differs from BGSS even when using parameterized expectations and simulation to compute moments. First, as we mentioned, we use a different center of expansion in the Taylor approximation of the value function. Second, we rely on recursion of the value function while BGSS rely on recursion of the policy function. Third, we use one-period-ahead simulations, while BGSS simulate paths from period 0 to the terminal period T .

⁸We use the term innovation in a generic sense and do not require δ_{t+1} to have zero mean for the purpose of flexibility. For instance, in the illustrative example of subsection 2.1, the innovation vector consists of the vector of risky asset returns \mathbf{R}_{t+1} and the exogenous state variable \mathbf{S}_{t+1} .

variable at time t depends on the choice variable at time $t - 1$. The structure of (18) is quite general: the function $u_t(\cdot)$ is a measure of immediate reward from taking action ω_t , and the function $\mathcal{H}_t(\cdot, \cdot)$ is an ‘‘aggregator’’ of the immediate reward $u_t(\mathbf{S}_t, \omega_t)$ and the expected value of future reward, $E_t[J_{t+1}(\mathbf{S}_{t+1})]$. Finite horizon problems with final date T are characterized by a known terminal condition satisfied by the value function $J_T(\mathbf{S}_T)$. Infinite horizon problems are also accommodated by the general recursion in (18), once we remove the explicit dependence of the value function from time, i.e., $J_t(\cdot) = J(\cdot)$ for all t .

To illustrate the generality of the setup described in (18) in the context of dynamic portfolio choice problems, note that, by taking an aggregator function $\mathcal{H}_t(\cdot, \cdot)$ of the form $\mathcal{H}_t(u, v) = u + \beta v$, $0 < \beta < 1$, the problem (18) simplifies to the familiar Bellman equation for the case of time separable utility, i.e.,

$$J_t(\mathbf{S}_t) = \max_{\omega_t} \{u_t(\mathbf{S}_t, \omega_t) + \beta E_t[J_{t+1}(\mathbf{S}_{t+1})]\}. \quad (19)$$

Alternatively, if we take $\mathcal{H}_t(u, v) = \left[(1 - \beta)u^{\frac{1}{\theta}} + \beta v^{\frac{1}{\theta}} \right]^{\theta}$, $\theta \neq 0$, we obtain the following recursion:

$$J_t(\mathbf{S}_t) = \max_{\omega_t} \left[(1 - \beta)u_t(\mathbf{S}_t, \omega_t)^{\frac{1}{\theta}} + \beta E_t[J_{t+1}(\mathbf{S}_{t+1})]^{\frac{1}{\theta}} \right]^{\theta}, \quad (20)$$

which is the Bellman equation characterizing the optimization problem of an agent with recursive preferences as in Epstein and Zin (1989, 1991).⁹

As discussed in the previous subsection, it is important, in terms of precision and computational efficiency, to work with a strictly monotonic transformation $V_t(\cdot)$ of the value function $J_t(\cdot)$, instead of $J_t(\cdot)$ itself. The appropriate transformation is typically suggested by the utility function u_t and the terminal condition. The transformed value function $V_t(\cdot)$, is implicitly determined by the relationship

$$J_t(\mathbf{S}_t) = \mathcal{U}(V_t(\mathbf{S}_t)) \quad (21)$$

where $\mathcal{U}(\cdot)$ is a strictly monotonic function.¹⁰ Using this transformation, the recursive equation (18) can be written in terms of the function $V_t(\cdot)$ as follows

$$\mathcal{U}(V_t(\mathbf{S}_t)) = \max_{\omega_t} \{ \mathcal{H}_t(u_t(\mathbf{S}_t, \omega_t), E_t[\mathcal{U}(V_{t+1}(\mathbf{S}_{t+1}))]) \}. \quad (22)$$

As illustrated in the previous subsection, the basic idea of the SVD method relies on the separation of the choice variables from the shocks to the state variables in a multiplicative fashion. This separation reduces the original problem into an approximate one that involves conditional expectations of quantities that do

⁹To see the equivalence between the formulation (20) and the original formulation in Epstein and Zin (1991), define $J \equiv V^{1-\gamma}$, $\gamma > 0$ and take $\theta = (1 - \gamma)/\rho$, $\rho < 1$, $\rho \neq 0$. Substituting in (20) delivers the Bellman equation (8) in Epstein and Zin (1991) with $\alpha = 1 - \gamma$ where γ is the constant of relative risk aversion, ρ the intertemporal substitution parameter and β a time preference parameter. Notice that when $\gamma = 1 - \rho$, we have $\theta = 1$ and (20) simplifies to the time-separable case (19).

¹⁰In the CRRA case of subsection 2.1, $\mathcal{U}(\cdot)$ is the utility of terminal wealth and V_t is the certainty equivalent function.

not depend on choice variables. We now describe in detail the three steps of the general SVD methodology. We assume that an approximation to the transformed value function $V_{t+1}(\cdot)$ has already been obtained at time $t + 1$, and show how to obtain the optimal allocation $\boldsymbol{\omega}_t$ and the transformed value function $V_t(\cdot)$ at time t .

Step 1. Decomposition of state variables

The first step consists of decomposing the innovation vector $\boldsymbol{\delta}_{t+1} = \boldsymbol{\delta}_{t+1}(\mathbf{S}_t, \boldsymbol{\varepsilon}_{t+1})$ in the Bellman equation (18) into a predictable component and the corresponding stochastic deviation:

$$\boldsymbol{\delta}_{t+1} = \mathbf{c}_{\delta,t} + \boldsymbol{\varepsilon}_{\delta,t+1} \quad (23)$$

where $\mathbf{c}_{\delta,t}$ is a function of the state variable \mathbf{S}_t , i. e., predictable, and the stochastic deviation $\boldsymbol{\varepsilon}_{\delta,t+1}$ is a function of the fundamental shock $\boldsymbol{\varepsilon}_{t+1}$. Using this decomposition, we can rewrite the law of motion of the state variables as

$$\mathbf{S}_{t+1} = \Gamma_t(\mathbf{S}_t, \boldsymbol{\omega}_t, \mathbf{c}_{\delta,t} + \boldsymbol{\varepsilon}_{\delta,t+1}). \quad (24)$$

In performing the above decomposition, one has the flexibility to choose whether to decompose only the innovation to the endogenous state variables, as we did in Step 1 of the illustrative example in Subsection 2.1, or both endogenous and exogenous state variables.

Step 2. Separation of choice variables from shocks

The goal of the second step is to achieve separation between the choice variables $\boldsymbol{\omega}_t$ and fundamental shocks $\boldsymbol{\varepsilon}_{t+1}$. This is achieved by using a Taylor expansion of the value function $\mathcal{U}(V_{t+1}(\mathbf{S}_{t+1}))$ with respect to the stochastic deviation $\boldsymbol{\varepsilon}_{\delta,t+1}$ centered at zero. As a result, we obtain an approximation that can be expressed as a *sum of products* of functions that are separable in the choice variable $\boldsymbol{\omega}_t$ and the stochastic deviation $\boldsymbol{\varepsilon}_{\delta,t+1}$ as follows:

$$\mathcal{U}(V_{t+1}(\mathbf{S}_{t+1})) \approx \sum_{m=1}^M A_{t+1,m}(\mathbf{S}_t, \boldsymbol{\omega}_t) \cdot B_{t+1,m}(\boldsymbol{\varepsilon}_{\delta,t+1}), \quad (25)$$

where the terms $A_{t+1,m}(\mathbf{S}_t, \boldsymbol{\omega}_t)$ involve partial derivatives of $\mathcal{U}(V_{t+1}(\Gamma_t(\mathbf{S}_t, \boldsymbol{\omega}_t, \mathbf{c}_{\delta,t} + \boldsymbol{\varepsilon}_{\delta,t+1})))$ with respect to $\boldsymbol{\varepsilon}_{\delta,t+1}$ evaluated at zero, and the terms $B_{t+1,m}(\boldsymbol{\varepsilon}_{\delta,t+1})$ are products of powers of elements of $\boldsymbol{\varepsilon}_{\delta,t+1}$. Note that the terms $A_{t+1,m}(\mathbf{S}_t, \boldsymbol{\omega}_t)$ depend on the choice variable $\boldsymbol{\omega}_t$ but do not depend on the fundamental shock $\boldsymbol{\varepsilon}_{t+1}$. In contrast, the terms $B_{t+1,m}(\boldsymbol{\varepsilon}_{\delta,t+1})$ only depend on the fundamental shock $\boldsymbol{\varepsilon}_{t+1}$. It follows from (25) that the conditional expectation $E_t[\mathcal{U}(V_{t+1}(\mathbf{S}_{t+1}))]$ in (22) can be approximated as

$$E_t[\mathcal{U}(V_{t+1}(\mathbf{S}_{t+1}))] \approx \sum_{m=1}^M A_{t+1,m}(\mathbf{S}_t, \boldsymbol{\omega}_t) \cdot E_t[B_{t+1,m}(\boldsymbol{\varepsilon}_{\delta,t+1})]. \quad (26)$$

The derivatives of the composite function $\mathcal{U}(V_{t+1}(\cdot))$ contained in the terms $A_{t+1,m}(\mathbf{S}_t, \boldsymbol{\omega}_t)$ can be efficiently computed by a recursive scheme known as the Faà di Bruno (1855, 1857) formula, as we discuss in Appendix A.1.

Step 3. Computation of conditional expectations

In the third and last step of our solution method, we compute the conditional expectations in (26). Note that, as a consequence of the previous two steps, the argument of the conditional expectation does not depend on the choice variables $\boldsymbol{\omega}_t$ and that the terms $B_{t+1,m}(\boldsymbol{\varepsilon}_{\delta,t+1})$ are products of powers of elements of $\boldsymbol{\varepsilon}_{\delta,t+1}$. The computation of these expectations can be carried out easily either analytically, under certain distributional assumptions, or numerically via quadrature or simulation methods. These possibilities result in three versions of the SVD approach that we label as SVD-A (analytical), SVD-Q (quadrature), and SVD-S (simulation).

We conclude this subsection with some further discussion on implementation issues. After computing the required conditional moments in step 3, we use the approximation (26) in (22) to obtain an approximate Bellman equation which yields an approximation to the optimal choice variable $\boldsymbol{\omega}_t$ and the transformed value function $V_t(\cdot)$ for each grid point in the state space. Next, we obtain an approximation to the transformed value function $V_t(\cdot)$ over the entire state space by a projection step and the procedure is repeated at time $t-1$. The method then proceeds backwards until time zero is reached. Different approaches are available for approximating the function $V_t(\cdot)$, such as discretization of the state space with interpolation or projection methods. The former is subject to the curse of dimensionality and, thus, highly inefficient for problems with multiple state variables. In our implementation, we rely on the projection method and, in each step of the backward induction, we approximate the value function by projecting it on a space of basis functions, such as polynomials or radial basis functions.¹¹ To evenly and efficiently cover the state space, we construct the state space grid using quasi-random (or low-discrepancy) sequences.¹² These sequences are multi-dimensional extensions of standard one-dimensional uniform grids that substantially enhance the computational efficiency of the SVD method.

As mentioned in the previous subsection, the choice of the center of expansion is critical for the convergence properties of the Taylor approximation. Intuitively, to contain the approximation error, one should choose the center of expansion to minimize the magnitude of the stochastic deviation $\boldsymbol{\varepsilon}_{\delta,t+1}$. Garlappi and Skoulakis (2008b) show theoretically that this goal is achieved by selecting the center (or midpoint) of the support of the innovation of the state variable $\boldsymbol{\delta}_{t+1}$ as the center of expansion $\mathbf{c}_{\delta,t}$.

¹¹See Haykin (1999) for a detailed discussion of radial basis functions.

¹²See Niederreiter (1992) for a comprehensive treatment of quasi-random sequences.

A final issue relates to the convergence of the Taylor series approximation. Although a formal treatment of this subject is beyond the scope of this paper, note that this issue does not pose a concern for the SVD method. In fact, one can guarantee convergence of the Taylor series approximation by appropriate selection of the transformation function \mathcal{U} . If \mathcal{U} does not have a globally convergent Taylor series representation, then one can approximate \mathcal{U} over the range of $V(\cdot)$ by a function that does, such as a polynomial or a radial basis function. Since the transformed value is approximated in the projection step by a polynomial or a radial basis function, it follows that the Taylor series approximation to $\mathcal{U}(V(\cdot))$ then converges for any center of expansion (see Theorem 9.25 in Apostol (1974)).

In the remainder of the paper, we apply the SVD methodology to solve a variety of static and dynamic portfolio choice problems.

3 Static portfolio choice problems

To illustrate the accuracy of the SVD approximation method, we first consider static portfolio choice problems with multiple risky assets. In subsection 3.1, we solve the portfolio choice problem of an investor with constant relative risk aversion (CRRA) preferences and normally distributed log excess returns. Because this problem does not admit a closed-form analytical solution, we use as a benchmark the approximate solution obtained by computing expectations via a quadrature method (e.g., Judd (1998)). In subsection 3.2, we argue that the specific form of Taylor expansion is critical for the accuracy of the approximation, and show that the SVD method, based on a decomposition of the log excess return, is unambiguously superior to the BGSS method, as documented by the certainty equivalent losses associated with the two methods.

3.1 CRRA preferences

Consider an investor with CRRA preferences represented by the utility function

$$U(W) = \frac{1}{1-\gamma} W^{1-\gamma}, \quad \gamma > 1. \quad (27)$$

The investment opportunity set consists of N risky assets and one risk-free asset. The portfolio choice problem is

$$\max_{\boldsymbol{\omega}} E[U(W_1)], \quad W_1 = W_0(R_f + \boldsymbol{\omega}'\mathbf{R}), \quad (28)$$

where $W_0 > 0$ denotes the current level of wealth, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_N)'$ is the N -dimensional vector of portfolio weights, and \mathbf{R} the N -dimensional vector of excess returns on the risky assets, i.e., the difference between gross return vector \mathbf{R}_g and the gross risk-free rate R_f , $\mathbf{R} = \mathbf{R}_g - R_f \mathbf{1}_N$. Since the problem is

homogeneous in wealth, we can assume, with no loss of generality, that $W_0 = 1$. To impose no short sales and borrowing constraints, we restrict the portfolio weights to be between 0% and 100%. We follow the portfolio choice literature (e.g., Campbell and Viceira (1999)) and assume that the vector of log excess returns, $\mathbf{r} = \log(\mathbf{R}_g) - \log(\mathbf{R}_f)\mathbf{1}_N$, is normally distributed, with mean $\boldsymbol{\mu}_r$ and covariance matrix $\boldsymbol{\Sigma}_r$, $\mathbf{r} \sim \mathcal{N}(\boldsymbol{\mu}_r, \boldsymbol{\Sigma}_r)$.

The SVD method can be applied in two alternative ways. The first employs a decomposition of the excess return \mathbf{R} , while the second employs a decomposition of the log excess return \mathbf{r} . It turns out the latter is much more efficient as it converges much faster. We next describe the two variants of the SVD method.

To apply the SVD method using the excess return decomposition, we write

$$\mathbf{R} = \mathbf{c}_R + \boldsymbol{\varepsilon}_R, \quad (29)$$

and use the M -th total order Taylor expansion of $U(W_1) = U(R_f + \boldsymbol{\omega}'(\mathbf{c}_R + \boldsymbol{\varepsilon}_R))$ with respect to $\boldsymbol{\varepsilon}_R$ centered at $\mathbf{0}_N$. This gives rise to the following approximate optimization problem

$$\max_{\boldsymbol{\omega}} \frac{1}{1-\gamma} \sum_{m=0}^M (1-\gamma)_m (R_f + \boldsymbol{\omega}'\mathbf{c}_R)^{1-\gamma-m} \sum_{\{\mathbf{m}:|\mathbf{m}|=m\}} \frac{1}{\mathbf{m}!} \cdot \prod_{n=1}^N \omega_n^{m_n} \cdot E \left[\prod_{n=1}^N \varepsilon_{R,n}^{m_n} \right], \quad (30)$$

where $(1-\gamma)_m = (1-\gamma)(-\gamma)\cdots(1-\gamma-m+1)$ is the m -th order falling factorial, and $|\mathbf{m}| = \sum_{n=1}^N m_n$, $\mathbf{m}! = \prod_{n=1}^N (m_n!)$ for any vector $\mathbf{m} = (m_1, \dots, m_N)$ of non-negative integers.

We choose $\mathbf{c}_R = (c_{R,1}, \dots, c_{R,N})'$ by setting $c_{R,n}$ equal to the *midpoint* of the support of R_n truncated at its 0.5 and 99.5 percentiles, for $n = 1, \dots, N$. See Garlappi and Skoulakis (2008b) for theoretical justification of this choice. Solving for the optimal portfolio weight vector $\boldsymbol{\omega}$ requires computing the cross moments $E \left[\prod_{n=1}^N \varepsilon_{R,n}^{m_n} \right]$ in (30). For normally distributed log excess returns, this computation can be performed in an efficient fashion using the recursive scheme of Skoulakis (2008). Under alternative distributional assumptions, the cross moments can be computed via quadrature or Monte-Carlo simulation methods.

To apply the SVD method using the log excess return decomposition, we write

$$\mathbf{r} = \mathbf{c}_r + \boldsymbol{\varepsilon}_r. \quad (31)$$

Since $\mathbf{R} = R_f(e^{\mathbf{r}} - \mathbf{1}_N)$, the portfolio return is given by $R_f(1 + \boldsymbol{\omega}'(e^{\mathbf{r}} - \mathbf{1}_N))$. It follows that $U(W_1) = R_f^{1-\gamma} h(\boldsymbol{\varepsilon}_r; \boldsymbol{\omega})$ where $h(\boldsymbol{\varepsilon}_r; \boldsymbol{\omega}) = \frac{1}{1-\gamma} (1 + \boldsymbol{\omega}'(e^{\mathbf{c}_r + \boldsymbol{\varepsilon}_r} - \mathbf{1}_N))^{1-\gamma}$ can be expressed as the composite function $h(\boldsymbol{\varepsilon}_r; \boldsymbol{\omega}) = f(g(\boldsymbol{\varepsilon}_r; \boldsymbol{\omega}))$ with $f(y) = \frac{y^{1-\gamma}}{1-\gamma}$, $y \in \mathbb{R}$ and $g(\boldsymbol{\varepsilon}_r; \boldsymbol{\omega}) = 1 + \boldsymbol{\omega}'(e^{\mathbf{c}_r + \boldsymbol{\varepsilon}_r} - \mathbf{1}_N)$, $\boldsymbol{\varepsilon}_r \in \mathbb{R}^N$. Using

a Taylor expansion with respect to $\boldsymbol{\varepsilon}_r$ centered at $\mathbf{0}_N$, we obtain the following approximate optimization problem

$$\max_{\boldsymbol{\omega}} \sum_{|\mathbf{m}| \leq M} h_{\mathbf{m}}(\boldsymbol{\omega}) E \left[\prod_{n=1}^N \varepsilon_{r,n}^{m_n} \right], \quad (32)$$

where $h_{\mathbf{m}}(\boldsymbol{\omega})$ denotes the \mathbf{m} -th order partial derivative of $h(\boldsymbol{\varepsilon}_r; \boldsymbol{\omega})$ with respect to $\boldsymbol{\varepsilon}_r$ evaluated at $\mathbf{0}_N$. We compute the partial derivatives of the composite function $h(\cdot) = f(g(\cdot))$ in an efficient fashion using a recursive scheme known as the Faà di Bruno formula as developed in Savits (2006). Since the distribution of $\boldsymbol{\varepsilon}_r$ is normal and hence symmetric, we select the \mathbf{c}_r to be the expected value of \mathbf{r} , i.e., $\mathbf{c}_r = \boldsymbol{\mu}_r$. Under this choice, $\boldsymbol{\varepsilon}_r$ follows a $\mathcal{N}(\mathbf{0}_N, \boldsymbol{\Sigma}_r)$ distribution and its cross moments are easily computed using the efficient recursive scheme of Savits (2006), as we explain in Appendix A.2.

We illustrate the SVD method by solving the portfolio choice problem of a CRRA investor with three risky assets. The three assets considered are the gross indices for USA, Europe and Pacific, obtained from Morgan Stanley Capital International (MSCI)-Barra. We use a time series from December 1969 to July 2006 to obtain parameter estimates and assume an annual risk-free rate of 5%. As the investment horizon increases, the distribution of returns shocks becomes wider, rendering the approximation less accurate. To check the validity of the approximation, we conservatively choose a horizon of one year. The estimated annual mean and covariance matrix of the log excess returns on the USA, Europe, and Pacific MSCI-Barra gross indices are:

$$\boldsymbol{\mu}_r = \begin{bmatrix} 0.0530 \\ 0.0620 \\ 0.0570 \end{bmatrix}, \quad \boldsymbol{\Sigma}_r = \begin{bmatrix} 0.0263 & & \\ 0.0219 & 0.0324 & \\ 0.0183 & 0.0282 & 0.0714 \end{bmatrix}. \quad (33)$$

We solve for the approximate optimal portfolio using these parameters and compare the SVD solution to that obtained by solving the optimization problem in (28) using numerical quadrature methods to approximate expectations, as in Judd (1998). Specifically, we use Gauss-Hermite quadrature with ten nodes in each dimension. We use this solution as a benchmark for comparison purposes.

Table 1 shows that, for CRRA preferences, the SVD method is extremely accurate, especially the version that uses decomposition of log excess return. The SVD portfolio weights are indistinguishable from their quadrature counterparts. Moreover, the associated certainty equivalent loss (CEL) with respect to the benchmark quadrature solution, stated in annualized basis points, is zero when the order of Taylor expansion, M , is as small as 4. While the SVD solution based on excess return decomposition does converge and eventually gives accurate results, it does so more slowly. See Garlappi and Skoulakis (2008b) for further discussion and evidence on the superior performance of the log excess return decomposition.

3.2 Comparison of the BGSS and SVD methods

We conclude this section by comparing the SVD method with the method proposed by BGSS. In the context of a static portfolio choice problem, the difference between the BGSS and the SVD approaches is the point around which the Taylor expansion of the utility function is performed. To see this difference clearly, consider the CRRA portfolio choice problem with normally distributed log excess returns studied in the previous subsection. The BGSS method uses a Taylor expansion of the portfolio return $R_f + \boldsymbol{\omega}'\mathbf{R}$ around the risk-free rate R_f . This is equivalent to a Taylor expansion with respect to the excess return \mathbf{R} around $\mathbf{0}_N$. The resulting approximate optimization problem in BGSS is

$$\max_{\boldsymbol{\omega}} \frac{1}{1-\gamma} \sum_{m=0}^M (1-\gamma)_m R_f^{1-\gamma-m} \sum_{\{\mathbf{m}:|\mathbf{m}|=m\}} \frac{1}{\mathbf{m}!} \cdot \prod_{n=1}^N \omega_n^{m_n} \cdot E \left[\prod_{n=1}^N R_n^{m_n} \right]. \quad (34)$$

BGSS suggest computing the expectations $E \left[\prod_{n=1}^N R_n^{m_n} \right]$ in the above expression via Monte-Carlo simulation. However, in the context of the static CRRA problem with normally distributed log excess return \mathbf{r} , these expectations can be computed analytically using the efficient recursive scheme in Skoulakis (2008).

In contrast, the SVD method uses a Taylor expansion with respect to the excess return \mathbf{R} around \mathbf{c}_R resulting in the approximate optimization problem described by equation (30). In our implementation, we choose \mathbf{c}_R to be the midpoint of the support of \mathbf{R} truncated at its 0.5 and 99.5 percentiles. An alternative version uses a Taylor expansion with respect to the log excess return \mathbf{r} centered at its mean. These two variants of the SVD method are discussed in detail in the previous subsection.

It turns out that the particular form of Taylor expansion is critical for the accuracy of the approximation. The BGSS expansion results in Taylor deviations of large magnitude and, hence, leads to significant approximation errors. In contrast, both variants of the SVD method ensure that the Taylor deviations are contained and the approximate solutions much more accurate. We refer the reader to Garlappi and Skoulakis (2008b) for a theoretical analysis of this issue. To illustrate the importance of using the appropriate Taylor approximation, we report in Table 2 a comparison between the BGSS and the SVD methods for the case of two risky assets. The log excess returns on the two risky assets are assumed to be jointly normally distributed with means equal to 8% and 11% and standard deviations equal to 13% and 20%, respectively. The correlation between the log excess returns on the two assets takes 5 values: -0.5 , -0.25 , 0 , 0.25 , and 0.5 . The annualized risk-free rate is set equal to 5% and the coefficient of relative risk aversion is set equal to $\gamma = 10$. We impose no short sales and borrowing constraints. The order of Taylor expansion ranges from 4 to 8 for all three methods. The table presents the certainty equivalent loss (CEL) in annualized basis points with respect to the benchmark solution obtained using Gauss-Hermite quadrature integration with 10 nodes in each dimension.

A few important points emerge from Table 2. First, the CEL associated with the BGSS solution can take extremely large values, as high as 1,500 basis points when an odd number of Taylor terms is used in the expansion. Second, the BGSS solutions oscillate wildly as the order M of the Taylor expansion alternates from even to odd. While striking, this result should not be surprising in light of previous evidence in the existing literature. Indeed, Loistl (1976) and Hlawitschka (1994) have pointed out this undesirable property of the Taylor approximation when the center of expansion is the mean of the risky asset return distribution. Third, the SVD method, by using the midpoint of the (truncated) support of the risky asset return distribution or an expansion with respect to the log excess return, ensures the stable convergence of the solution. Finally, we observe that, while the SVD variant that uses decomposition of the excess return does converge, the version that uses decomposition of the log excess return converges much faster. See Garlappi and Skoulakis (2008b) for a theoretical justification of the good convergence properties of the log excess return decomposition and Garlappi and Skoulakis (2008a) for further evidence on the accuracy of this method in the context of a dynamic portfolio choice with CRRA preferences and predictable risky asset returns.

In summary, the evidence in this section illustrates that the SVD method can be extremely accurate in solving static portfolio choice problems. In the rest of the paper, we extend our analysis to dynamic portfolio choice problems with both constant and time-varying investment opportunity sets.

4 Dynamic portfolio choice problems

We begin the analysis of the SVD method in a dynamic context by considering the problem of an investor with CARA preferences and a constant investment opportunity set with i.i.d. normally distributed risky asset returns. This relatively simple setting allows us to achieve two goals. First, since under these conditions the dynamic portfolio choice problem admits a closed-form solution, we can assess the accuracy of the SVD method against an exact benchmark (Subsection 4.1). Second, this setup also facilitates a comparison of the approximations provided by the SVD and BGSS methods in a dynamic setting (Subsection 4.2). In Section 5, we extend our analysis to the case of a time-varying investment opportunity set and solve for the optimal consumption and portfolio choices over the life cycle of an investor with recursive preferences facing predictable asset returns.

4.1 Accuracy of the SVD method

Consider an investor with CARA preferences who seeks to maximize the expected utility of terminal wealth, $E_0 [U(W_T)]$, where $U(W_T) = -\exp(-\alpha W_T)$ and α is the coefficient of absolute risk aversion. Let R_f denote

the risk-free rate and \mathbf{R}_t denote the vector of excess returns on the risky assets at time t . The investment opportunity set is constant, and the excess returns \mathbf{R}_t are assumed to be i.i.d. normally distributed with mean $\boldsymbol{\mu}_R$ and covariance matrix $\boldsymbol{\Sigma}_R$, i.e., $\mathbf{R}_t \sim \mathcal{N}(\boldsymbol{\mu}_R, \boldsymbol{\Sigma}_R)$ for all t . Since returns are i.i.d., the only state variable in the problem is the wealth W_t . The budget equation is

$$W_{t+1} = W_t(R_f + \boldsymbol{\omega}'_t \mathbf{R}_{t+1}), \quad t = 0, \dots, T-1, \quad (35)$$

where $\boldsymbol{\omega}_t$ denotes the risky asset portfolio weights. Let $J_t(W_t)$ be the value function which satisfies the usual Bellman equation, $J_t(W_t) = \max_{\boldsymbol{\omega}_t} E_t [J_{t+1}(W_t(R_f + \boldsymbol{\omega}'_t \mathbf{R}_{t+1}))]$, with terminal condition $J_T(W_T) = -\exp(-\alpha W_T)$. The following proposition provides the analytical solution to the portfolio choice problem.¹³

Proposition 1 *Assume that the investment opportunity set is characterized by risky assets whose returns \mathbf{R}_t in excess of the risk free rate R_f are i.i.d. normally distributed with mean $\boldsymbol{\mu}_R$ and covariance matrix $\boldsymbol{\Sigma}_R$, i.e., $\mathbf{R}_t \sim \mathcal{N}(\boldsymbol{\mu}_R, \boldsymbol{\Sigma}_R)$. Then, the value function $J_t(W_t)$ and the optimal portfolio $\boldsymbol{\omega}_t^*(W_t)$ chosen by an investor with CARA preferences $U(W) = -e^{-\alpha W}$, $\alpha \neq 0$, and investment horizon T , are given by the following expressions*

$$J_t(W_t) = -\exp\left(-\alpha W_t R_f^{T-t} - \frac{T-t}{2} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R\right), \quad t = 0, \dots, T \quad (36)$$

$$\boldsymbol{\omega}_t^*(W_t) = \frac{1}{\alpha W_t R_f^{(T-1)-t}} \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R, \quad t = 0, \dots, T-1. \quad (37)$$

The closed-form solution in the above proposition allows us to assess the accuracy of the SVD approach in solving the dynamic problem of a CARA investor.

As discussed in Section 2, instead of using directly the value function $J_t(\cdot)$, it is computationally more efficient to work with the *certainty equivalent* $V_t(W_t)$ of $J_t(W_t)$, defined by the relationship $J_t(W_t) = U(V_t(W_t))$ or,

$$V_t(W_t) = -\frac{1}{\alpha} \log(-J_t(W_t)). \quad (38)$$

Using this transformation, the dynamic program becomes $U(V_t(W_t)) = \max_{\boldsymbol{\omega}_t} E_t[U(V_{t+1}(W_{t+1}))]$ or, from the definition of $U(\cdot)$,

$$V_t(W_t) = -\frac{1}{\alpha} \log\left(-\max_{\boldsymbol{\omega}_t} E_t\left[-e^{-\alpha V_{t+1}(W_{t+1})}\right]\right) \quad (39)$$

with terminal condition $V_T(W_T) = W_T$. We solve the problem using backward induction as follows. We first decompose the risky asset excess return \mathbf{R}_{t+1} into its mean $\boldsymbol{\mu}_R$ and the associated deviation $\boldsymbol{\varepsilon}_{R,t+1}$, i.e., $\mathbf{R}_{t+1} = \boldsymbol{\mu}_R + \boldsymbol{\varepsilon}_{R,t+1}$ where $\boldsymbol{\varepsilon}_{R,t+1} \sim \mathcal{N}(\mathbf{0}_N, \boldsymbol{\Sigma}_R)$.¹⁴ The evolution of the wealth process is then described

¹³The proof of the proposition is a standard application of backward induction and is available upon request.

¹⁴Since the distribution of asset returns is symmetric, the choice of the mean as center of expansion is desirable as it minimizes the largest possible deviation in the Taylor approximation. See Garlappi and Skoulakis (2008b) for details.

by

$$W_{t+1} = W_t (R_f + \boldsymbol{\omega}'_t \mathbf{R}_{t+1}) = \mu_{W,t} + \varepsilon_{W,t+1} \quad (40)$$

where $\mu_{W,t} = W_t(R_f + \boldsymbol{\omega}'_t \boldsymbol{\mu}_R)$ and $\varepsilon_{W,t+1} = W_t(\boldsymbol{\omega}'_t \boldsymbol{\varepsilon}_{R,t+1})$. Suppose that at time $t + 1$ we have an approximation of the certainty equivalent function $V_{t+1}(W_{t+1})$ as a polynomial of order K_{t+1} in W_{t+1} :

$$V_{t+1}(W_{t+1}) \approx \sum_{k=0}^{K_{t+1}} \vartheta_{t+1,k} W_{t+1}^k = \sum_{k=0}^{K_{t+1}} \vartheta_{t+1,k} (\mu_{W,t} + \varepsilon_{W,t+1})^k. \quad (41)$$

Note that at time T , the certainty equivalent function is $V_T(W_T) = W_T$, and therefore the above approximation is exact for $\vartheta_{T,0} = 0$, $\vartheta_{T,1} = 1$, and $\vartheta_{T,k} = 0$ for $k \geq 2$. Using the decomposition (40) and the approximation (41), the conditional expectation in the Bellman equation (39) is approximated as follows

$$E_t \left[-e^{-\alpha V_{t+1}(W_{t+1})} \right] \approx E_t \left[-e^{-\alpha \sum_{k=0}^{K_{t+1}} \vartheta_{t+1,k} (\mu_{W,t} + \varepsilon_{W,t+1})^k} \right] \equiv E_t [g_{t+1}(\varepsilon_{W,t+1})]. \quad (42)$$

Following the SVD approach, we approximate the function g_{t+1} by a sum of products of quantities that depend only on the portfolio weights $\boldsymbol{\omega}_t$ and quantities that depend only on the return shocks $\boldsymbol{\varepsilon}_{R,t+1}$. To achieve such a goal, we first use a Taylor expansion of order M_{t+1} centered at $\varepsilon_{W,t+1} = 0$:

$$g_{t+1}(\varepsilon_{W,t+1}) \approx \sum_{m=0}^{M_{t+1}} \frac{1}{m!} g_{t+1}^{(m)}(0) \varepsilon_{W,t+1}^m, \quad (43)$$

where $g_{t+1}^{(m)}(0)$ denotes the m -th derivative of $g_{t+1}(\varepsilon_{W,t+1})$ evaluated at $\varepsilon_{W,t+1} = 0$, for $m = 0, \dots, M_{t+1}$. Second, using the fact that $\varepsilon_{W,t} = \boldsymbol{\omega}'_t \boldsymbol{\varepsilon}_{R,t+1}$, we fully separate the choice variable $\boldsymbol{\omega}_t$ from $\boldsymbol{\varepsilon}_{R,t+1}$, by applying the multinomial formula to the expression $\varepsilon_{W,t+1}^m$, and obtain

$$\varepsilon_{W,t+1}^m = (\boldsymbol{\omega}'_t \boldsymbol{\varepsilon}_{R,t+1})^m = \sum_{\{\mathbf{q}: |\mathbf{q}|=m\}} \frac{m!}{\mathbf{q}!} \prod_{n=1}^N \omega_{t,n}^{q_n} \prod_{n=1}^N \varepsilon_{R_n,t+1}^{q_n}, \quad (44)$$

where $\mathbf{q} = (q_1, \dots, q_N)$ is a vector of nonnegative integers, $|\mathbf{q}| = q_1 + \dots + q_N$, and $\mathbf{q}! = q_1! \dots q_N!$. Combining (39), (42), (43), and (44), we can compute an approximation to the certainty equivalent function at time t as follows

$$V_t(W_t) \approx -\frac{1}{\alpha} \log \left(-\max_{\boldsymbol{\omega}_t} \left\{ \sum_{m=0}^{M_{t+1}} g_{t+1}^{(m)}(0) \sum_{\{\mathbf{q}: |\mathbf{q}|=m\}} \frac{1}{\mathbf{q}!} \prod_{n=1}^N \omega_{t,n}^{q_n} E_t \left[\prod_{n=1}^N \varepsilon_{R_n,t+1}^{q_n} \right] \right\} \right). \quad (45)$$

Solving the optimization problem in (45) involves computing the moments $E_t \left[\prod_{n=1}^N \varepsilon_{R_n,t+1}^{q_n} \right]$. This calculation can be performed in an efficient fashion using the recursive Faà di Bruno scheme developed in Savits (2006). See Appendix A.2 for details.

We construct a grid on the wealth space,¹⁵ solve the Bellman equation (45) at each point on the grid, and obtain an approximation to $V_t(W_t)$ by fitting a polynomial of degree K_t to this solution, as discussed above. The backward induction then proceeds in the standard fashion until we reach time zero.

To highlight the performance of the SVD method, we compare its approximate solution to the exact closed-form solution provided in Proposition 1. Our calibration is based on the international data that we used in Subsection 3.1. The estimated annual mean and covariance matrix of the excess returns on the USA, Europe, and Pacific MSCI-Barra gross indices are¹⁶

$$\boldsymbol{\mu}_R = \begin{bmatrix} 0.0712 \\ 0.0854 \\ 0.1023 \end{bmatrix}, \quad \boldsymbol{\Sigma}_R = \begin{bmatrix} 0.0292 & & \\ 0.0251 & 0.0427 & \\ 0.0190 & 0.0347 & 0.0999 \end{bmatrix}. \quad (46)$$

The annual risk-free rate is set equal to 5%. We consider three levels of absolute risk aversion $\alpha = 2, 4,$ and $6,$ and three choices for the investment horizon $T = 10, 20,$ and 30 years. We approximate the certainty equivalent function V_t by a quadratic function of wealth, i.e., $K_{t+1} = 2$ in (41) for all $t,$ and use a Taylor approximation of order $M_{t+1} = 4$ for all $t.$ Table 3 presents the certainty equivalent returns (CER) that correspond to the exact solution provided by Proposition 1 and the SVD solution for five levels of initial wealth $W_0 = 1, 1.25, 1.5, 1.75,$ and $2.$ The CER represents the annualized risk-free return that the investor is willing to accept in exchange of the opportunity to invest optimally in the existing risky assets over the next T years. To define CER formally, consider a portfolio policy denoted by $\hat{\omega}_t(\cdot), t = 0, \dots, T - 1.$ The associated terminal wealth \hat{W}_T is obtained by the recursion $\hat{W}_{t+1} = \hat{W}_t \left(1 + \hat{\omega}_t(\hat{W}_t)' \mathbf{R}_{t+1} \right), t = 0, \dots, T - 1,$ with $\hat{W}_0 = W_0.$ The annualized CER is then defined through the equation $U(W_0(1 + CER)^T) = E_0[U(\hat{W}_T)].$ For the exact solution described in Proposition 1, we have $E_0[U(\hat{W}_T)] = J_0(W_0)$ and, therefore, we can use equation (37) to explicitly express the CER associated with the exact solution as

$$CER_{\text{EX}} = \left(\frac{1}{W_0} U^{-1}(J_0(W_0)) \right)^{\frac{1}{T}} - 1 = \left(R_f^T + \frac{T}{2\alpha W_0} \boldsymbol{\mu}_R \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R \right)^{\frac{1}{T}} - 1. \quad (47)$$

Since the value function associated with the SVD method is not available analytically, we compute the CER for the SVD solution through Monte Carlo simulation. Specifically, we simulate I risky asset return paths $\{\mathbf{R}_{t+1}^i : t = 0, \dots, T - 1\},$ for $i = 1, \dots, I,$ according to the i.i.d. data generating process

¹⁵More precisely, starting at time 0 and for each period, we construct a sequence of expanding grids ensuring that all attainable next-period wealth levels, obtained using a set of feasible portfolio allocations, lie within the limits of the next-period wealth grid with sufficiently high probability. We increase the number of points on the grid and the coverage probability until we obtain a stable solution. In this particular implementation, the results were obtained using, in each period, 100 grid points and coverage probability of 95%.

¹⁶Note that the parameter estimates in (46) refer to excess returns, while the parameter estimates in (33), used in Subsection 3.1, refer to *log* excess returns.

$\mathbf{R}_{t+1} \sim \mathcal{N}(\boldsymbol{\mu}_R, \boldsymbol{\Sigma}_R)$. For the i -th simulated path, we obtain the associated terminal wealth $\hat{W}_{T,i}$ and estimate the expected utility of terminal wealth by

$$EU_{\text{SVD}} = \frac{1}{I} \sum_{i=1}^I U(\hat{W}_{T,i}). \quad (48)$$

The estimated CER associated with the SVD solution is then obtained by

$$CER_{\text{SVD}} = \left(\frac{1}{W_0} U^{-1}(EU_{\text{SVD}}) \right)^{\frac{1}{T}} - 1. \quad (49)$$

In our implementation we use one million return paths with antithetic random numbers,¹⁷ in order to reduce the Monte Carlo error.

From Table 3, which reports the CER in percentages, it is clear that the SVD approximation is extremely accurate. The difference in CER is less than one basis point in all cases considered. Note that frequently the CER obtained by the SVD solution is higher (by a fraction of a basis point) than the CER obtained by the exact solution. This discrepancy is purely due to Monte Carlo error, since, by construction, the CER obtained by the approximate solution cannot be larger than the CER obtained by the exact solution.

4.2 Comparison of the BGSS and SVD methods

The analytically tractable framework of this section allows us to investigate in depth the differences between the SVD and BGSS approaches in the context of dynamic portfolio problems. As discussed in Section 3, in the case of static portfolio choice problems, the two approaches differ only in the choice of the center of expansion used to approximate the terminal utility. The SVD method allows for flexibility in such a choice, which turns out to be crucial for the accuracy of the approximation, as illustrated in Section 3.

In the case of dynamic portfolio choice problems, however, the two methods differ in another important dimension. While SVD is a *value function iteration* method, BGSS is a *policy iteration* method. Specifically, for each time t , SVD uses the next-period value function $V_{t+1}(\cdot)$ to determine the current-period optimal portfolio weights $\boldsymbol{\omega}_t^*(\cdot)$ and the value function V_t , as it moves backwards from time T to time 0. In contrast, to determine the optimal portfolio weights at time t , BGSS relies on the optimal portfolio weights for the periods $t+1$ through T , that have already been computed in the previous steps of the backward recursion. The conditional expectations involved in the approximate Bellman equation are computed through cross-sectional regressions based on simulated paths of asset returns and state variables.

The simulation-based policy-iteration procedure in BGSS avoids the need of recursively approximating the value function, as is required by a value function iteration approach like the SVD. However, there are

¹⁷See Glasserman (2004) for details on variance reduction using antithetic random numbers.

two important difficulties associated with such an approach. First, simulating state variables forward is practically feasible only when the variables to be simulated are exogenous. If a state variable is endogenous, e.g., wealth in the CARA problem of this section, one needs to know the choice variable before simulating the endogenous state variable. But the choice variable is also the ultimate objective of the optimization. This circularity makes the problem of simulating an endogenous state variable cumbersome and practically not feasible.

Second, when approximating the expected utility of terminal wealth using a Taylor expansion, one needs to make sure to take the potential dependence of optimal portfolio weights on wealth explicitly into account. By the very nature of the policy iteration approach, the BGSS method is forced to ignore the dependence of future optimal portfolio weights on current wealth. This is evident in equations (16) and (17) on p. 843 in BGSS where future optimal weights are treated as independent of current wealth. Such a property holds for portfolio choice problems with homothetic preferences and no frictions, but is generally violated in all other cases, as in the CARA example considered in this section.

To highlight the consequences of ignoring the dependence of portfolio weights on wealth, we compare the BGSS and SVD approximations to the exact solution in the context of a simple two-period problem with CARA preferences. This framework provides also a clean way of investigating the convergence properties of both methods, as the number of terms in the Taylor expansion increases.

Let us consider the portfolio choice problem with CARA preferences and i.i.d. normally distributed returns studied in the previous subsection over two periods, i.e., $T = 2$. We assume that the optimal portfolio at time 1, $\omega_1^*(W_1)$, has been determined and the only task left is to determine the portfolio at time zero. Both SVD and BGSS methods rely on an approximate Bellman equation of the following form

$$J_0(W_0) = \max_{\omega_0} E_0[J_1(W_1)] \approx \max_{\omega_0} E_0 \left[\sum_{m=0}^M \frac{1}{m!} \mathcal{D}^m J_1(W_1)|_{W_1=\bar{W}_1} (W_1 - \bar{W}_1)^m \right], \quad (50)$$

where $\mathcal{D}^m J_1(W_1) = \frac{d^m J_1(W_1)}{dW_1^m}$, $m = 0, 1, \dots, M$. The approximate Bellman equation is obtained by the M -th order Taylor expansion of $J_1(W_1)$ centered at \bar{W}_1 .

In order to provide a fair and transparent comparison between SVD and BGSS, we assume that at time zero, both methods make use of the *exact* optimal portfolio weight $\omega_1^*(W_1)$, or equivalently the *exact* value function $J_1(W_1)$, at time 1. For future reference, note that

$$\omega_1^*(W_1) = \frac{1}{\alpha W_1} \Sigma_R^{-1} \mu_R, \quad (51)$$

and

$$J_1(W_1) = -\exp \left(-\alpha W_1 R_f - \frac{1}{2} \mu_R' \Sigma_R^{-1} \mu_R \right), \quad (52)$$

as it follows from Proposition 1.

There are two important differences between the SVD and BGSS methods. The first difference concerns the center of Taylor expansion \bar{W}_1 . As discussed in Subsection 4.1, SVD chooses $\bar{W}_1 = \mu_{W,1} = W_0(R_f + \omega'_0 \mu_R)$ and, hence, $W_1 - \bar{W}_1 = W_0(\omega'_0 \varepsilon_{R,1}) = \varepsilon_{W,1}$, where $\varepsilon_{R,1} = \mathbf{R}_1 - \mu_R$. In contrast, BGSS chooses $\bar{W}_1 = W_0 R_f$ which implies $W_1 - \bar{W}_1 = W_0(\omega'_0 R_1)$. The second difference between the two methods concerns the computation of the derivatives of the value function $\mathcal{D}^m J_1(W_1) = \frac{d^m J_1(W_1)}{dW_1^m}$. The SVD method, which is based on value function iteration, explicitly uses the value function $J_1(W_1)$ when computing these derivatives. In contrast, the BGSS method, which does not use value function iteration, relies on the law of motion

$$W_2 = W_1(R_f + \omega_1^*(W_1)' \mathbf{R}_2) \quad (53)$$

to express the value function $J_1(W_1)$ as

$$J_1(W_1) = E_1 [U(W_2)] = E_1 [U(W_1(R_f + \omega_1^*(W_1)' \mathbf{R}_2))]. \quad (54)$$

However, when computing derivatives, the BGSS method ignores the dependence of $\omega_1^*(W_1)$ on W_1 (see equations (16) and (17) on p. 843 in BGSS). Therefore, the m -th order derivative in (50), for $m = 0, 1, \dots, M$, evaluated at $W_1 = W_0 R_f$, is computed by the BGSS method as

$$\mathcal{D}_{\text{BGSS}}^m J_1(W_1)|_{W_1=W_0 R_f} = E_1 \left[U^{(m)} \left(\bar{W}_2^{\text{BGSS}} \right) (R_f + \omega_1^*(W_0 R_f)' \mathbf{R}_2)^m \right], \quad (55)$$

where

$$\bar{W}_2^{\text{BGSS}} = W_0 R_f (R_f + \omega_1^*(W_0 R_f)' \mathbf{R}_2). \quad (56)$$

Using the expression for the optimal portfolio weight $\omega_1^*(W_1)$ given in (51), we obtain

$$\bar{W}_2^{\text{BGSS}} = W_0 R_f^2 + \frac{1}{\alpha} \mu'_R \Sigma_R^{-1} \mathbf{R}_2. \quad (57)$$

and therefore

$$\mathcal{D}_{\text{BGSS}}^m J_1(W_1)|_{W_1=W_0 R_f} = E_1 \left[U^{(m)} \left(\bar{W}_2^{\text{BGSS}} \right) \left(R_f + \frac{1}{\alpha W_0 R_f} \mu'_R \Sigma_R^{-1} \mathbf{R}_2 \right)^m \right]. \quad (58)$$

The following lemma provides an explicit expression for the above derivative.

Lemma 1 *Let \widetilde{W}_1 be any value of wealth at time 1 and define $\widetilde{W}_2 = \widetilde{W}_1 R_f + \frac{1}{\alpha} \mu'_R \Sigma_R^{-1} \mathbf{R}_2$. Then,*

$$E_1 \left[U^{(m)} \left(\widetilde{W}_2 \right) \left(R_f + \frac{1}{\alpha \widetilde{W}_1} \mu'_R \Sigma_R^{-1} \mathbf{R}_2 \right)^m \right] = (-1)^{m+1} \alpha^m e^{-\alpha \widetilde{W}_1 R_f - \frac{s^2}{2}} \sum_{k=0}^m \binom{m}{k} \left(\frac{1}{\alpha \widetilde{W}_1} \right)^k R_f^{m-k} S^k \phi_k, \quad (59)$$

where $S^2 = \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R$ denotes the squared Sharpe ratio, and ϕ_k denotes the k -th order central moment of the standard normal distribution given by

$$\phi_k = \begin{cases} 0, & \text{if } k \text{ odd} \\ \frac{k!}{(k/2)!2^{k/2}}, & \text{if } k \text{ even.} \end{cases} \quad (60)$$

Applying Lemma 1 with $\widetilde{W}_1 = W_0 R_f$, we obtain

$$D_{\text{BGSS}}^m \equiv \mathcal{D}_{\text{BGSS}}^m J_1(W_1)|_{W_1=W_0 R_f} = (-1)^{m+1} \alpha^m e^{-\alpha W_0 R_f^2 - \frac{S^2}{2}} \sum_{k=0}^m \binom{m}{k} \left(\frac{1}{\alpha W_0 R_f} \right)^k S^k R_f^{m-k} \phi_k. \quad (61)$$

Therefore, by (i) assuming a Taylor expansion centered at $\overline{W}_1 = W_0 R_f$ and (ii) ignoring the dependence of the optimal portfolio weight $\boldsymbol{\omega}_1^*(W_1)$ on wealth W_1 when computing derivatives, the BGSS method solves the following approximate Bellman equation

$$J_0(W_0) \approx \max_{\boldsymbol{\omega}_0} E_0 \left[\sum_{m=0}^M \frac{1}{m!} D_{\text{BGSS}}^m (W_0 \boldsymbol{\omega}'_0 \mathbf{R}_1)^m \right] \quad (62)$$

$$= \max_{\boldsymbol{\omega}_0} \left[\sum_{m=0}^M D_{\text{BGSS}}^m W_0^m \sum_{\{\mathbf{q}:|\mathbf{q}|=m\}} \frac{1}{\mathbf{q}!} \prod_{n=1}^N \omega_{0,n}^{q_n} E_0 \left[\prod_{n=1}^N R_{1,n}^{q_n} \right] \right], \quad (63)$$

where D_{BGSS}^m , $m = 0, 1, \dots, M$, are given by (61).

To explicitly obtain the approximate Bellman equation used by the SVD method, we need an expression for the derivatives of the value function $J_1(W_1)$ evaluated at $\mu_{W,1} = W_0(R_f + \boldsymbol{\omega}'_0 \boldsymbol{\mu}_R)$. It follows from (52) that, for $m = 0, 1, \dots, M$, we have

$$\frac{d^m J_1(W_1)}{dW_1^m} = (-1)^{m+1} \alpha^m R_f^m e^{-\alpha W_1 R_f - \frac{S^2}{2}}. \quad (64)$$

Hence, the derivatives of $J_1(W_1)$ evaluated at $\mu_{W,1}$, as used by the SVD method, are

$$D_{\text{SVD}}^m(\boldsymbol{\omega}_0) \equiv \mathcal{D}_{\text{SVD}}^m J_1(W_1)|_{W_1=\mu_{W,1}} = \frac{d^m J_1(\mu_{W,1})}{dW_1^m} = (-1)^{m+1} \alpha^m R_f^m e^{-\alpha W_0(R_f + \boldsymbol{\omega}'_0 \boldsymbol{\mu}_R) R_f - \frac{S^2}{2}}. \quad (65)$$

Therefore, by (i) assuming a Taylor expansion centered at $\overline{W}_1 = \mu_{W,1} \equiv W_0(R_f + \boldsymbol{\omega}'_0 \boldsymbol{\mu})$ and (ii) explicitly using the value function $J_1(W_1)$, the SVD method solves the following approximate Bellman equation

$$J_0(W_0) \approx \max_{\boldsymbol{\omega}_0} E_0 \left[\sum_{m=0}^M \frac{1}{m!} D_{\text{SVD}}^m(\boldsymbol{\omega}_0) \varepsilon_{W,1}^m \right] \quad (66)$$

$$= \max_{\boldsymbol{\omega}_0} \left[\sum_{m=0}^M D_{\text{SVD}}^m(\boldsymbol{\omega}_0) W_0^m \sum_{\{\mathbf{q}:|\mathbf{q}|=m\}} \frac{1}{\mathbf{q}!} \prod_{n=1}^N \omega_{0,n}^{q_n} E_0 \left[\prod_{n=1}^N \varepsilon_{1,n}^{q_n} \right] \right], \quad (67)$$

where the coefficients $D_{\text{SVD}}^m(\boldsymbol{\omega}_0)$, $m = 0, 1, \dots, M$, are given by (65).

The expectations $E_0 \left[\prod_{n=1}^N R_{1,n}^{q_n} \right]$ and $E_0 \left[\prod_{n=1}^N \varepsilon_{1,n}^{q_n} \right]$ for all $\mathbf{q} = (q_1, \dots, q_N)$ with $|\mathbf{q}| \leq M$, that appear in the approximate Bellman equations (62)-(63) and (66)-(67), can be efficiently computed using the recursive scheme of Savits (2006) as illustrated in Appendix A.2.

To better understand the two dimensions along which SVD and BGSS differ, we develop two modifications that can be thought of as ‘‘hybrids’’ of the BGSS and SVD approaches. In the first modification, referred to as M1, we (i) use $\bar{W}_1 = \mu_{W,1} = W_0(R_f + \boldsymbol{\omega}'_0 \boldsymbol{\mu}_R)$ as the center of expansion (as does the SVD approach), and (ii) ignore the dependence of $\boldsymbol{\omega}_1^*(W_1)$ on W_1 when computing the derivatives with respect to W_1 (as does the BGSS approach). In the second modification, referred to as M2, we (i) use $W_0 R_f$ as the center of expansion (as does the BGSS approach), and (ii) explicitly use the value function $J_1(W_1)$ in the computation of derivatives (as does the SVD approach).

It follows that the approximate Bellman equation used by modification M1 is identical to the SVD approximate Bellman equation (66)-(67) except that the derivatives of $J_1(W_1)$ are computed differently. Specifically, since the M1 method ignores the dependence of $\boldsymbol{\omega}_1^*(W_1)$ on W_1 , the m -th order derivative in (50), for $m = 0, 1, \dots, M$, evaluated at $W_1 = \mu_{W,1}$, is computed by the M1 method as

$$\mathcal{D}_{\text{M1}}^m J_1(W_1)|_{W_1=\mu_{W,1}} = E_1 \left[U^{(m)} \left(\bar{W}_2^{\text{M1}} \right) \left(R_f + \boldsymbol{\omega}_1^*(\mu_{W,1})' \mathbf{R}_2 \right)^m \right], \quad (68)$$

where

$$\bar{W}_2^{\text{M1}} = \mu_{W,1} (R_f + \boldsymbol{\omega}_1^*(\mu_{W,1})' \mathbf{R}_2). \quad (69)$$

Using the expression for the optimal portfolio weight $\boldsymbol{\omega}_1^*(W_1)$ given in (51) yields

$$\bar{W}_2^{\text{M1}} = \mu_{W,1} R_f + \frac{1}{\alpha} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \mathbf{R}_2 \quad (70)$$

and therefore

$$\mathcal{D}_{\text{M1}}^m J_1(W_1)|_{W_1=\mu_{W,1}} = E_1 \left[U^{(m)} \left(\bar{W}_2^{\text{M1}} \right) \left(R_f + \frac{1}{\alpha \mu_{W,1}} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \mathbf{R}_2 \right)^m \right]. \quad (71)$$

Applying Lemma 1 with $\tilde{W}_1 = \mu_{W,1}$, we obtain

$$D_{\text{M1}}^m(\boldsymbol{\omega}_0) \equiv \mathcal{D}_{\text{M1}}^m J_1(W_1)|_{W_1=\mu_{W,1}} = (-1)^{m+1} \alpha^m e^{-\alpha \mu_{W,1} R_f - \frac{\alpha^2}{2}} \sum_{k=0}^m \binom{m}{k} \left(\frac{1}{\alpha \mu_{W,1}} \right)^k S^k R_f^{m-k} \phi_k. \quad (72)$$

Hence, the approximate Bellman equation solved by the M1 method is

$$\begin{aligned} J_0(W_0) &\approx \max_{\boldsymbol{\omega}_0} E_0 \left[\sum_{m=0}^M \frac{1}{m!} D_{\text{M1}}^m(\boldsymbol{\omega}_0) \varepsilon_{W,1}^m \right] \\ &= \max_{\boldsymbol{\omega}_0} \left[\sum_{m=0}^M D_{\text{M1}}^m(\boldsymbol{\omega}_0) W_0^m \sum_{\{\mathbf{q}:|\mathbf{q}|=m\}} \frac{1}{\mathbf{q}!} \prod_{n=1}^N \omega_{0,n}^{q_n} E_0 \left[\prod_{n=1}^N \varepsilon_{1,n}^{q_n} \right] \right], \end{aligned} \quad (73)$$

where the coefficients $D_{M1}^m(\boldsymbol{\omega}_0)$, $m = 0, 1, \dots, M$, are given by (72). Note that $D_{M1}^m(\boldsymbol{\omega}_0)$ is explicitly dependent on the portfolio weight $\boldsymbol{\omega}_0$.

The approximate Bellman equation used by modification M2 has the same structure as the BGSS approximate Bellman equation (62)-(63) except that the derivatives of $J_1(W_1)$ are computed differently. Specifically, since the M2 method explicitly uses the value function $J_1(W_1)$, the derivatives of $J_1(W_1)$ evaluated at $W_0 R_f$, as used by M2, are given by

$$D_{M2}^m \equiv \mathcal{D}_{M2}^m J_1(W_1)|_{W_1=W_0 R_f} = \frac{d^m J_1(W_0 R_f)}{dW_1^m} = (-1)^{m+1} \alpha^m R_f^m e^{-\alpha W_0 R_f^2 - \frac{s^2}{2}}. \quad (74)$$

as it follows from (64). Formally, this leads to the following approximate Bellman equation solved by the M2 method

$$J_0(W_0) \approx \max_{\boldsymbol{\omega}_0} E_0 \left[\sum_{m=0}^M \frac{1}{m!} D_{M2}^m (W_0 \boldsymbol{\omega}'_0 \mathbf{R}_1)^m \right] \quad (75)$$

$$= \max_{\boldsymbol{\omega}_0} \left[\sum_{m=0}^M D_{M2}^m W_0^m \sum_{\{\mathbf{q}:|\mathbf{q}|=m\}} \frac{1}{\mathbf{q}!} \prod_{n=1}^N \omega_{0,n}^{q_n} E_0 \left[\prod_{n=1}^N R_{1,n}^{q_n} \right] \right]. \quad (76)$$

where the coefficients D_{M2}^m , which no longer depend on the portfolio weights $\boldsymbol{\omega}_0$, are given by (74).

It is important to stress that accurate computation of derivatives is possible in methods which rely on the information contained in the value function, such as the SVD and M2 methods. In contrast, methods that rely on policy iteration and subsequently ignore dependence of portfolio weights on wealth, such as the BGSS and M1 methods, result in erroneous computation of derivatives and, therefore, in inaccurate solutions to the portfolio choice problem. Overall, the SVD method that uses value function iteration and appropriate center of expansion emerges as a dominant method, as we illustrate next.

We use each of the aforementioned four methods (BGSS, M1, M2, and SVD) to derive approximate optimal portfolio allocations $\boldsymbol{\omega}_0$ and compare the various solutions on the basis of their certainty equivalent return. Note that all methods make use of the exact optimal portfolio allocation $\boldsymbol{\omega}_1^*(W_1)$ and, therefore, the time-2 final wealth can be expressed as a function of the initial allocation $\boldsymbol{\omega}_0$ and the returns \mathbf{R}_1 and \mathbf{R}_2 . Specifically, using (51) we obtain

$$W_2 = W_1 (R_f + \boldsymbol{\omega}_1^*(W_1)' \mathbf{R}_2) = W_1 R_f + \frac{1}{\alpha} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \mathbf{R}_2 \quad (77)$$

and since $W_1 = W_0 (R_f + \boldsymbol{\omega}'_0 \mathbf{R}_1)$ we have

$$W_2 = W_0 (R_f + \boldsymbol{\omega}'_0 \mathbf{R}_1) R_f + \frac{1}{\alpha} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \mathbf{R}_2. \quad (78)$$

Hence, the expected utility of terminal wealth for initial wealth W_0 and portfolio allocation $\boldsymbol{\omega}_0$ can be computed as follows

$$\begin{aligned}
\mathcal{J}(W_0, \boldsymbol{\omega}_0) &= E_0 \left[-e^{-\alpha W_2} \right] \\
&= E_0 \left[-e^{-\alpha (W_0(R_f + \boldsymbol{\omega}'_0 \mathbf{R}_1)R_f + \frac{1}{\alpha} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \mathbf{R}_2)} \right] \\
&= -e^{-\alpha W_0(R_f + \boldsymbol{\omega}'_0 \boldsymbol{\mu}_R)R_f + \frac{1}{2} \alpha^2 W_0^2 R_f^2 (\boldsymbol{\omega}'_0 \boldsymbol{\Sigma}_R \boldsymbol{\omega}_0) - \frac{1}{2} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R}.
\end{aligned} \tag{79}$$

The last equality uses the independence of \mathbf{R}_1 and \mathbf{R}_2 and the formula for the moment generating function of the normal distribution.

Since we focus on a two-period problem, the per-period certainty equivalent return associated with the allocation $\boldsymbol{\omega}_0$ for initial wealth W_0 , denoted by $CER(W_0, \boldsymbol{\omega}_0)$, is implicitly defined by

$$\mathcal{J}_0(W_0, \boldsymbol{\omega}_0) = U \left(W_0(1 + CER(W_0, \boldsymbol{\omega}_0))^2 \right) = -e^{-\alpha W_0(1 + CER(W_0, \boldsymbol{\omega}_0))^2}. \tag{80}$$

Equivalently, we have $CER(W_0, \boldsymbol{\omega}_0) = \left[\frac{1}{W_0} U^{-1}(\mathcal{J}_0(W_0, \boldsymbol{\omega}_0)) \right]^{\frac{1}{2}} - 1$, or more explicitly

$$CER(W_0, \boldsymbol{\omega}_0) = \left[\frac{1}{W_0} \left(W_0(R_f + \boldsymbol{\omega}'_0 \boldsymbol{\mu}_R)R_f - \frac{1}{2} \alpha W_0^2 R_f^2 (\boldsymbol{\omega}'_0 \boldsymbol{\Sigma}_R \boldsymbol{\omega}_0) + \frac{1}{2\alpha} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R \right) \right]^{\frac{1}{2}} - 1. \tag{81}$$

The CER associated with the exact optimal solution is obtained from (47) for $T = 2$:

$$CER_{\text{EX}}(W_0) = \left[R_f^2 + \frac{1}{\alpha W_0} \boldsymbol{\mu}'_R \boldsymbol{\Sigma}_R^{-1} \boldsymbol{\mu}_R \right]^{\frac{1}{2}} - 1. \tag{82}$$

We assess the quality of each approximation by computing the certainty equivalent loss (CEL) of each method, defined as the difference between the CER obtained from the exact optimal solution given in Proposition 1 and the CER obtained from each approximate optimal solution.

Table 4 reports the CEL for the four methods described above. The numbers in the table refer to a problem with two risky assets and a risk free asset. The return on the risk free asset is set to 5% per annum, the annual expected excess returns of the risky assets are 10% and 15%, and their annual volatilities are 15% and 25%, respectively. Returns are assumed to be normally distributed and i.i.d. over time. The coefficient of absolute risk aversion is set equal to $\alpha = 4$, and the initial wealth level, W_0 , takes 3 values: 0.5, 1, and 1.5. The table reports the CEL (in basis points per annum) of the four solution methods for different values of (i) the correlation (ρ) across assets varies and (ii) the order of expansion M in the Taylor approximation.

The certainty equivalent losses reported in Table 4 imply that the SVD method is clearly superior to the other three methods, for all levels of asset return correlation and initial wealth considered. In

particular, using a Taylor expansion of order $M = 6$ or higher the SVD method produces a solution that is virtually indistinguishable from the exact solution, for all levels of initial wealth and asset return correlation. In contrast, with a Taylor expansions of order $M = 8$, the BGSS method accumulates losses between 0.16 ($W_0 = 1.5$, $\rho = 0.25$) and 107.67 basis points ($W_0 = 0.5$, $\rho = -0.25$). The choice of center of expansion seems to be important also for dynamic problems, confirming our findings in the solution of the static CRRA problem in Section 3. For example, for the BGSS and M2 methods, which choose $W_0 R_f$ as expansion point, the CEL are 239.48 and 105.22 basis points, respectively, when $\rho = -0.25$ and the order of expansion is $M = 4$. For the M2 and SVD approaches, which choose the expected next-period wealth $\mu_{W,1}$ as expansion point, the corresponding CEL are 2.57 and 21.96 basis points, respectively.

All methods, with the exception of M1, improve as the order of expansion M increases. The M1 method exhibits the rather undesirable property that its performance deteriorates as more terms used in the Taylor expansion. This is not surprising since this method does not properly account for the dependence of portfolio weights on wealth and the problem becomes more severe as higher-order derivatives are used (i.e., higher M). The good performance of the M1 method for $M = 4$ appears to be attributable to the good choice of the expansion point. However, the overall evidence suggests that the M1 method should not be relied upon in the solution of larger scale problems.

In summary, the analysis in this section illustrates that methods that correctly account for the dependence of portfolio weights on wealth (SVD and M2) are much more accurate than methods that ignore such dependence (BGSS and M1). We deduce that the use of Taylor expansions for solving approximate portfolio choice problems in which wealth is not a redundant state variable cannot ignore the value function and its properly computed derivatives. By choosing the expected wealth $\mu_{W,1}$ as center of expansion and computing derivatives of the value function properly, the SVD method emerges as the dominant method for addressing this class of problems.

In the next section, we relax the assumption of a time-invariant opportunity set and exploit the full potential of the SVD method to address a realistic life-cycle strategic asset allocation problem in which the investment opportunity set is governed by a large number of state variables.

5 Strategic asset allocation with predictable returns

We study the intertemporal consumption and portfolio choice problem of a finitely-lived investor with recursive preferences (à la Epstein and Zin (1989, 1991)), who can invest in multiple assets and faces a time-varying investment opportunity set with predictable risk premia of asset returns. It is well-known that

time variation in the investment opportunity set introduces hedging motives in the formation of optimal portfolios. Our setup is similar to that in Campbell, Chan, and Viceira (2003) (CCV hereafter) with three important differences: (i) the investor has a finite horizon instead of being infinitely-lived, (ii) the investor faces realistic borrowing and short-selling constraints, and (iii) our approximation scheme relies on decomposing the state variables instead of log-linearizing the budget constraint.

The finite horizon and the presence of portfolio constraints make the problem more realistic but, at the same time, considerably more challenging because the log-linearization methodology of CCV cannot be used for the solution of a constrained problem. Campbell, Cocco, Gomes, Maenhout, and Viceira (2001) use standard quadrature techniques to numerically solve the constrained version of the problem solved by Campbell and Viceira (1999), with one risky asset and one state variable. In our application we consider, as in CCV, *three* assets and *six* state variables. The presence of constraints is dictated also by the necessity to avoid bankruptcy in the portfolio problem. In the unconstrained problem solved by CCV, bankruptcy can occur with positive probability, in which case the utility function is not well-defined. Finally, our methodology allows us to consider a wide array of parameter specifications. This is particularly important for parameters like the elasticity of intertemporal substitution (EIS). It is well-known that estimates of the elasticity of intertemporal substitution vary substantially depending on whether one uses aggregate consumption data or household data and on the class of investors chosen (stock-holders or bond-holders). Vissing-Jørgensen (2002), for example, estimates EIS values of 0.3–0.4 for stock-holders and 0.8–1 for bond-holders. While the log-linear approximation relies on values of the EIS close to unity, the SVD method is accurate for *any* value of this parameter.

In the rest of this section, we describe the ingredients of the intertemporal portfolio problem, we show how to apply the SVD methodology in this setting, and illustrate the precision of the SVD method by comparing its results to those obtained from applying traditional quadrature methods.

5.1 A life-cycle portfolio problem with predictable returns

We consider an investor who allocates resources across three risky assets: nominal Treasury bills (the short-term asset), stocks, and long term nominal Treasury bonds.¹⁸ The predictability of asset risk premia is captured via a vector autoregressive (VAR) system that includes six state variables: the returns on the three assets (i.e., the short-term ex-post real interest rate, the excess stock return, and the excess bond return) plus three other variables that that empirical research has identified as useful return predictors

¹⁸Note that, because of inflation risk, the *real* return on the nominal bill is risky and hence the short-term asset is not a “risk-free” asset.

(i.e., the short-term nominal interest rate, the dividend-price ratio and the yield spread between long-term and short term bonds).

We denote by R_f^N the gross nominal return on the T-bill rate, by R_b^N the gross nominal return on the bond, by R_s^N the gross nominal stock return, and by Π the gross rate of inflation. The corresponding *real* gross returns on the three asset are: $R_f = R_f^N/\Pi$, $R_b = R_b^N/\Pi$ and $R_s = R_s^N/\Pi$, and the continuously compounded real returns are

$$r_f = \log(R_f), \quad r_b = \log(R_b), \quad r_s = \log(R_s). \quad (83)$$

Let $r_1 \equiv r_f$, $r_2 \equiv r_b - r_f$, $r_3 \equiv r_s - r_f$ and denote by \mathbf{r} the three-dimensional vector composed of r_f and the log excess returns on the bond and on the stock, i.e., $\mathbf{r} = (r_1, r_2, r_3)'$. The real return $R_p(\boldsymbol{\omega})$ on a given portfolio with weights $\boldsymbol{\omega} = (\omega_f, \omega_b, \omega_s)$ in the T-bill, T-bond, and stock, respectively, is

$$\begin{aligned} R_p(\boldsymbol{\omega}) &= R_f + \omega_b(R_b - R_f) + \omega_s(R_s - R_f) \\ &= e^{r_1} [1 + \omega_b(e^{r_2} - 1) + \omega_s(e^{r_3} - 1)], \end{aligned} \quad (84)$$

where the last equality follows from (83). Because the investor is not allowed to short any of the assets, each of the portfolio weights has to lie in the unit interval, i.e., $0 \leq \omega_i \leq 1$, $i = f, s, b$.

We collect the three variables used as predictors of asset returns in the vector $\mathbf{z} = (z_1, z_2, z_3)'$, where z_1 is the yield on the 90-day T-bill, z_2 is the dividend-price ratio, and z_3 is the spread between the 5-year zero-coupon bond yield and the T-bill rate. The six-dimensional state variable $\mathbf{y}_t = (\mathbf{r}'_t, \mathbf{z}'_t)'$ follows the VAR system

$$\mathbf{y}_{t+1} = \mathbf{a} + \mathbf{B}\mathbf{y}_t + \boldsymbol{\varepsilon}_{y,t+1}, \quad (85)$$

where \mathbf{a} is a 6×1 vector of intercepts, \mathbf{B} is a 6×6 matrix of slope coefficients and $\boldsymbol{\varepsilon}_{y,t+1}$ are the shocks to the state variable, satisfying the following distributional assumptions:

$$\boldsymbol{\varepsilon}_{y,t+1} \sim \mathcal{N}(\mathbf{0}_6, \boldsymbol{\Sigma}_\varepsilon), \quad \boldsymbol{\Sigma}_\varepsilon \equiv \text{Var}(\boldsymbol{\varepsilon}_{y,t+1}) = \begin{bmatrix} \boldsymbol{\Sigma}_{rr} & \boldsymbol{\Sigma}'_{rz} \\ \boldsymbol{\Sigma}_{rz} & \boldsymbol{\Sigma}_{zz} \end{bmatrix}. \quad (86)$$

The VAR specification (85) summarizes the dependence of asset returns on their lagged realizations as well as the other predictive variables. The distributional assumptions in (86) imply that the shocks are correlated cross-sectionally but not in the time series dimension. Note that if the covariance matrix $\boldsymbol{\Sigma}_\varepsilon$ has full rank, then shocks to the variables governing the evolution of the investment opportunity set are imperfectly correlated to shocks to asset return and therefore cannot be hedged completely by trading in these assets. In other words, in this setting, markets are incomplete.

The finitely-lived investor has recursive preferences of the form described by Epstein and Zin (1989). Specifically, the utility U_t at time t is defined recursively by

$$U_t = \left\{ (1 - \beta)C_t^\rho + \beta \left(E_t \left[U_{t+1}^{1-\gamma} \right] \right)^{\frac{\rho}{1-\gamma}} \right\}^{\frac{1}{\rho}}, \quad \rho \leq 1, \quad \rho \neq 0, \quad \beta > 0, \quad (87)$$

where C_t is consumption at time t , and the utility at the terminal date T is given by $U_T = (1 - \beta)^{1/\rho} W_T$ with W_T denoting terminal wealth. The investor finances consumption entirely from financial wealth and does not receive labor income. Defining by W_t the wealth at time t , the intertemporal budget constraint hence is

$$W_{t+1} = (W_t - C_t)R_{p,t+1}(\boldsymbol{\omega}_t). \quad (88)$$

Under these conditions, Epstein and Zin (1989) show that the Bellman equation of the portfolio choice problem takes the form

$$V_t(W_t, \mathbf{y}_t) = \max_{C_t, \boldsymbol{\omega}_t} \left\{ (1 - \beta)C_t^\rho + \beta \left(E_t \left(V_{t+1}^{1-\gamma}(W_{t+1}, \mathbf{y}_{t+1}) \right) \right)^{\frac{\rho}{1-\gamma}} \right\}^{\frac{1}{\rho}}, \quad (89)$$

where $\mathbf{y}_{t+1} = (\mathbf{r}'_{t+1}, \mathbf{z}'_{t+1})'$ is the vector of exogenous state variables defined above, γ is the coefficient of relative risk aversion, and $1/(1 - \rho)$ the elasticity of intertemporal substitution. In the time-separable case, risk aversion equals the inverse of the elasticity of intertemporal substitution, i.e., $\gamma = 1 - \rho$. In this case (89) reduces to the familiar Bellman equation for the consumption-portfolio choice problem in Merton (1973). Note finally that, by construction, the function V_t in the Bellman equation (89) can be interpreted as the certainty equivalent function discussed in Section 2. We next elaborate on the steps of the SVD methodology in the context of the problem under consideration.

5.2 Applying the SVD methodology

In principle, one could apply the SVD method directly to the Bellman equation (89). However, the homotheticity of preferences allows us to simplify the problem by removing wealth as a state variable according to the following lemma which restates Proposition 1 in Bhamra and Uppal (2006).

Lemma 2 *Under the homothetic recursive preferences described by (87), the value function that solves (89) is given by*

$$V_t(W_t, \mathbf{y}_t) = (1 - \beta)^{1/\rho} \mathcal{V}_t(\mathbf{y}_t) W_t, \quad (90)$$

where

$$\mathcal{V}_t(\mathbf{y}_t) = \left\{ 1 + \left[\beta \left(\min_{\boldsymbol{\omega}_t} E_t \left[(R_{p,t+1}(\boldsymbol{\omega}_t))^{1-\gamma} \mathcal{V}_{t+1}(\mathbf{y}_{t+1})^{1-\gamma} \right] \right)^{\frac{\rho}{1-\gamma}} \right]^{\frac{1}{1-\rho}} \right\}^{\frac{1-\rho}{\rho}}, \quad (91)$$

with terminal condition $\mathcal{V}_T(\mathbf{y}_T) = 1$. The optimal consumption-wealth ratio $c_t = C_t/W_t$ is given by

$$c_t = \mathcal{V}_t(\mathbf{y}_t)^{-\frac{\rho}{1-\rho}}. \quad (92)$$

The above lemma shows that the portfolio and consumption problem are separable. For a given value of the state variable \mathbf{y}_t , the optimal portfolio can then be found by solving the optimization problem

$$\min_{\boldsymbol{\omega}_t} E_t \left[(R_{p,t+1}(\boldsymbol{\omega}_t))^{1-\gamma} \mathcal{V}_{t+1}(\mathbf{y}_{t+1})^{1-\gamma} \right]. \quad (93)$$

The conditional expectations involved in the transformed Bellman equation (91) have the same form as the ones encountered in the previous sections. Therefore, the SVD methodology can be easily applied to the case of recursive preferences.

To apply the SVD method, we first decompose each state variable $y_{i,t+1}$ in the vector $\mathbf{y}_{t+1} = (\mathbf{r}'_{t+1}, \mathbf{z}'_{t+1})'$ into its conditional mean $\mu_{i,t}$ and the associated stochastic zero-mean innovation $\varepsilon_{i,t+1}$ as follows

$$y_{i,t+1} = \mu_{i,t} + \varepsilon_{i,t+1}, \quad i = 1, \dots, 6. \quad (94)$$

Substituting the decomposition (94) in equation (93), we can then express the quantities $R_{p,t+1}(\boldsymbol{\omega}_t)^{1-\gamma}$ and $\mathcal{V}_{t+1}(\mathbf{y}_{t+1})^{1-\gamma}$ in (93) as functions of the innovations $\boldsymbol{\varepsilon}_{\mathbf{y},t+1}$. We do so by introducing two functions h_1 and h_2 defined as follows

$$h_1(\boldsymbol{\omega}_t, \varepsilon_{1,t+1}, \varepsilon_{2,t+1}, \varepsilon_{3,t+1}) = \left(e^{\mu_{1,t} + \varepsilon_{1,t+1}} [1 + \omega_{b,t}(e^{\mu_{2,t} + \varepsilon_{2,t+1}} - 1) + \omega_{s,t}(e^{\mu_{3,t} + \varepsilon_{3,t+1}} - 1)] \right)^{1-\gamma} \quad (95)$$

$$h_2(\boldsymbol{\varepsilon}_{\mathbf{y},t+1}) = \left(\mathcal{V}_{t+1}(\boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_{\mathbf{y},t+1}) \right)^{1-\gamma}, \quad (96)$$

where the base of the power on the right-hand side of (95) follows directly from the definition of $R_{p,t+1}(\boldsymbol{\omega}_t)$ in (84) and the decomposition (94). Having defined these two functions, we can express the expectation in (93) as

$$E_t [R_{p,t+1}(\boldsymbol{\omega}_t)^{1-\gamma} \cdot \mathcal{V}_{t+1}(\mathbf{y}_{t+1})^{1-\gamma}] = E_t [h_1(\boldsymbol{\omega}_t, \varepsilon_{1,t+1}, \varepsilon_{2,t+1}, \varepsilon_{3,t+1}) \cdot h_2(\boldsymbol{\varepsilon}_{\mathbf{y},t+1})]. \quad (97)$$

In order to facilitate the efficient computation of the above expectation, the SVD method proceeds by employing a Taylor expansion of the functions h_1 and h_2 in (97) with respect to $(\varepsilon_{1,t+1}, \varepsilon_{2,t+1}, \varepsilon_{3,t+1})$ and $\boldsymbol{\varepsilon}_{\mathbf{y},t+1}$ around $\mathbf{0}_3$ and $\mathbf{0}_6$, respectively.

Because h_1 is a function of 3 state variables only (the asset returns) while h_2 is a function of all 6 state variables (3 asset returns and 3 exogenous predictors), to simplify notation while performing Taylor expansions of these functions, we denote by \mathbf{n} the vector of indices of the first three variables, i.e., $\mathbf{n} = (n_1, n_2, n_3)$, and by \mathbf{k} the vector of indices for all six variables, i.e., $\mathbf{k} = (k_1, \dots, k_6)$. These indices

represents the order of partial derivative in each of the corresponding variables. We further denote by $h_{1,\mathbf{n}}$ the partial derivative of order \mathbf{n} of the function h_1 , i.e.,

$$h_{1,\mathbf{n}}(\boldsymbol{\omega}_t) = \frac{\partial^{|\mathbf{n}|} h_1}{\partial \varepsilon_1^{n_1} \partial \varepsilon_2^{n_2} \partial \varepsilon_3^{n_3}}(\boldsymbol{\omega}_t, \mathbf{0}_3), \quad (98)$$

and by $h_{2,\mathbf{k}}$ the partial derivative of order \mathbf{k} of the function h_2 , i.e.,

$$h_{2,\mathbf{k}} = \frac{\partial^{|\mathbf{k}|} h_2}{\partial \varepsilon_1^{k_1} \partial \varepsilon_2^{k_2} \partial \varepsilon_3^{k_3} \partial \varepsilon_4^{k_4} \partial \varepsilon_5^{k_5} \partial \varepsilon_6^{k_6}}(\mathbf{0}_6). \quad (99)$$

The expectation (97) is then approximated by

$$E_t \left[\left(\sum_{|\mathbf{n}| \leq M_1} \frac{1}{\mathbf{n}!} h_{1,\mathbf{n}}(\boldsymbol{\omega}_t) \prod_{i=1}^3 \varepsilon_{i,t+1}^{n_i} \right) \cdot \left(\sum_{|\mathbf{k}| \leq M_2} \frac{1}{\mathbf{k}!} h_{2,\mathbf{k}} \prod_{i=1}^6 \varepsilon_{i,t+1}^{k_i} \right) \right], \quad (100)$$

where M_1 and M_2 are the chosen orders of Taylor expansion for h_1 and h_2 , respectively. The scalars $h_{1,\mathbf{n}}(\boldsymbol{\omega}_t)$ and $h_{2,\mathbf{k}}$ can be efficiently computed recursively, using the results in Savits (2006) (see Appendix A.1). It follows from (100) that we are left with the task of computing the cross moments of the form:

$$E_t \left[\prod_{i=1}^3 \varepsilon_{i,t+1}^{n_i} \prod_{i=1}^6 \varepsilon_{i,t+1}^{k_i} \right] = E_t \left[\prod_{i=1}^3 \varepsilon_{i,t+1}^{n_i+k_i} \prod_{i=4}^6 \varepsilon_{i,t+1}^{k_i} \right]. \quad (101)$$

Under the VAR dynamics (85), this amounts to computing the central moments of a multivariate normal distribution for which the efficient recursive schemes of Savits (2006) (see Appendix A.2) can be used. Notice that the computation of these moments needs to be done *only once* and does not have to be repeated at each step of the backward recursion. In a high-dimensional problem like the one considered in this section, this provides a tremendous improvement in computational efficiency as we document shortly.

In our implementation of the SVD method, we use Taylor expansions of order 4 for the functions h_1 and h_2 (i.e., $M_1 = M_2 = 4$ in (100)). The problem is solved on a sequence of successive grids in the six-dimensional state space, starting from the stationary distribution of the state variable vector \mathbf{y}_t implied by the VAR specification, and going forward ensuring that subsequent realizations of the state variable are within the next-period grid with a probability of at least 99%. In each period, the grid consists of 1,000 evenly distributed points.¹⁹ At every step of the backward induction, we approximate the value function $\mathcal{V}_t(\mathbf{y}_t)$ by a radial basis function with 500 Gaussian kernels.²⁰ Given the high dimensionality and

¹⁹We select the grid points using quasi-random (or low-discrepancy) sequences which are multi-dimensional extensions of standard one-dimensional uniform grids. Niederreiter (1992) provides a thorough treatment of such sequences.

²⁰A radial basis function (RBF) is a linear combination of a number of Gaussian kernels. Specifically, the functional form of an RBF defined on \mathbb{R}^d is given by

$$\mathcal{V}(\mathbf{y}) = a + \sum_{k=1}^K b_k \cdot e^{-\theta_k^2 \|\mathbf{y} - \mathbf{c}_k\|}$$

where $a, b_1, \dots, b_K, \theta_1, \dots, \theta_K$ are real numbers and $\mathbf{c}_1, \dots, \mathbf{c}_K$ (centers) are d -dimensional vectors. See Haykin (1999) for further details on radial basis functions. Numerical methods for approximating functions using RBF are readily available.

the complexity of the value function approximation problem, we opted for the more powerful class of radial basis functions instead of using more traditional basis functions such as polynomials. The use of radial basis functions allows us to keep the mean square error in the approximation of $\mathcal{V}_t(\cdot)$ in each step of the backward recursion in the order of 10^{-6} .

5.3 Results

The data, at the annual frequency, and the estimation of the VAR system (85) are taken directly from CCV.²¹ The data span a period from 1890 to 1998. The choice of an annual frequency is dictated by the exigence of creating a more challenging environment to assess the quality of the SVD approximation. In our benchmark calibration of the recursive utility specification (89), we set the coefficient of risk aversion γ to 5, the elasticity of intertemporal substitution $1/(1 - \rho)$ to 0.5, and the time preference parameter β to 0.94.

We assess the accuracy of the SVD method by comparing its solution to the benchmark solution obtained via a quadrature method with 4 quadrature points in each of the six dimensions of the state space. It is worth mentioning that quadrature methods are extremely inefficient for problems of this scale.²²

Table 5 displays the optimal portfolio allocations for a planning horizon of 15 years (Panel A) and 30 years (Panel B). The solution of the SVD method is reported in the columns labeled ‘SVD’ and the solution from the benchmark quadrature method is in the column labeled ‘Q’. Each of the 6 subpanels in the table reports the optimal portfolio allocation to T-bonds (ω_b) and stocks (ω_s), and the optimal consumption-to-wealth ratios (c).²³ To assess the accuracy of the solution, for each of the 6 predictor variables, we record the optimal portfolio weights and consumption at the first (p_{25}) second (p_{50}) and third (p_{75}) quartile of the stationary distribution of that predictor, *conditional* on keeping all the remaining 5 predictors at their median level. For example, the first subpanel reports the optimal portfolio allocation and consumption policy at the 25th, 50th and 75th percentile of the distribution of the log return on the nominal T-bill, assuming that the other 5 predictors (excess nominal T-bond log return, excess stock log return, short-term rate, dividend yield and yield spread) are set at their median values. Because the quartiles are computed according to the stationary distributions, the portfolio weights at different time horizons in each table are directly comparable since they refer to the same point in the state space.

²¹We thank John Campbell for making the data available. Sample statistics and parameter estimates can be found in Tables 1 and 2 of CCV.

²²For example, in our MATLAB implementation, the SVD method took 3.46 hours while the quadrature method took 126.59 hours which translates into a huge improvement in terms of efficiency by a factor of at least 35.

²³The portfolio holdings ω_f of T-bills are $\omega_f = 1 - \omega_b - \omega_s$.

A comparison between the entries in the SVD- and Q-columns reveals that the SVD method is extremely accurate, producing asset allocation weights and consumption policies that are never more than 1% apart (and in the vast majority of cases less than 0.5% apart) from the corresponding quantities obtained using the quadrature method. These results illustrate that, while much more computationally efficient than the quadrature method, the SVD method is quite precise and reliable.

Since stocks have a large and positive Sharpe ratio (CCV estimate it to be 0.374 in their Table 1), the optimal portfolio is usually long in equity. Moreover, because positive shocks to the dividend yield are positively correlated with stock expected return and tend to be associated with lower realized stock return, equity represent a natural hedge against changes in the opportunity set driven by the dividend yield. As a consequence the optimal allocation to stocks increases with the dividend yield for both investment horizons (15 and 30 years), as illustrated in the ‘Dividend yield’ subpanel of Table 5.

Focusing attention on the three macro predictors we note that the dividend yield and the yield spread affect significantly the holdings in the T-bond and stock. Not surprisingly, holdings of stocks are relatively more affected by the dividend yield while holdings of T-bond react more to changes in the yield spread. The difference of stock holding between the first and the third quartile is about 35%, independently of the investment horizon. In contrast, the portfolio holdings do not seem to be affected in a substantial way by changes in the short term nominal rate. Finally, there does not seem to be an appreciable difference between portfolio weights of investors with different horizons, although in general the holdings of T-bond tend to decline and those of equity tend to increase as the horizon lengthens.

To understand the impact of different values of the elasticity of intertemporal substitution (EIS), we employ the SVD method to solve the problem for two different values of EIS: 0.75 and 1.25. In Table 6, we report the asset allocations to T-bond and stock as well as the consumption-to-wealth ratio. Because EIS captures the willingness of the investor to substitute consumption intertemporally, the quantity that is most affected by changes in EIS is the consumption-to-wealth ratio. To understand the pattern of consumption over the state space, it is useful to recall that changes in expected returns on the asset carry both *income* and *substitution* effects. If a shock to one or more state variables causes the expected portfolio return to increase, then the investor will trade off the incentive to consume more out of his wealth (income effect) with the incentive to save and take advantage of the favorable investment opportunities (substitution effect). For EIS smaller than one the income effects dominates while for EIS higher than one the substitution effect dominates. We can see this in Table 6 by looking, for example, at the behavior of the consumption-to-wealth ratio at different level of the dividend-yield. As dividend yield increases, for

low EIS the consumption-to-wealth ratio increases, while for high EIS the ratio decreases. All else being equal the effect of a change in EIS is more pronounced for longer investment horizons.

Finally, in Table 7 we decompose the total demand in the three assets into *myopic* (M) and *hedging* (H) demand. We define myopic demand as the portfolio choice of an investor who solves a static problem, i.e. he ignores the evolution of the state variables describing the investment opportunity set. As the magnitude of the hedging demands illustrates, ignoring the evolution of the investment opportunity set leads to significantly different portfolios. The magnitude of the hedging demands are smaller than those reported in CCV because we are imposing short-selling and borrowing constraints. The hedging demand for stock is always positive and increasing with the investment horizon. A positive hedging demand in a particular asset emerges when the shock to the return on that asset is negatively correlated to the shock to state variable that most affect the holding in that asset. For stocks this variable is the dividend yield and positive holdings of stocks are a way to hedge against changes in the evolution of the opportunity set brought forth by such variable. However, with six state variables as in this case, there are multiple cross effects that render the interpretation of each hedging demand less straightforward. The no-short selling constraint appears to be more severe in the T-bill. This is consistent with the results in CCV who find that at the annual level, the hedging demand for bond and stock is positive, thus causing a large short position in cash. Because our investor cannot borrow, he optimally holds a zero position in the T-bill over a large portion of the state space, as confirmed by the values of ω_f in Table 7.

6 Conclusion

We propose a new, precise, and efficient numerical method for the solution of dynamic portfolio choice problems and, more generally, for a broad class of dynamic programming problems. The key advantage of the State Variable Decomposition (SVD) method we develop is to approximate the conditional expectations of the value function as polynomials in suitably chosen exogenous random shocks to the state variable. This is achieved via (i) a decomposition of each state variable into a predictable component and the associated stochastic deviations and (ii) the use of Taylor approximation. It follows then that the choice variables can be “factored out” from the computation of conditional expectations, simplifying the problem of computing conditional expectation to one of computing conditional moments of random variables. In several cases, such as when shocks follow normal or log-normal distributions, these moments are available in closed form.

In applying the method, we take particular care in minimizing the sources of approximation errors that are associated with our numerical scheme. In particular, to minimize the projection error we rely on approximation to the certainty equivalent function instead of the original value function. To minimize the

error in the use of Taylor approximation, we select the predictable part of the state variable decomposition to guarantee convergence of the Taylor series. This allows us to control the approximation error by adding terms in the Taylor expansion.

We apply the SVD method to a broad array of static and dynamic problems and verify its accuracy under several realistic features such as the presence of intermediate consumption, multiple risky assets and state variables, non-additive preferences, time-varying investment opportunity sets, and non-redundant endogenous state variables. In all of the applications we consider, the SVD method emerges as a fast, accurate, and reliable procedure making it a flexible and important new toolkit for solving dynamic problems in finance and economics.

A Useful recursive schemes from multivariate analysis

In this appendix, we summarize two useful recursive schemes that facilitate the efficient numerical implementation of the SVD approach. Appendix A.1 introduces a recursive scheme for the computation of derivatives of a composite function of multiple variables and Appendix A.2 describes a recursive scheme for the efficient computation of central moments of a multivariate normal distribution. The results in this appendix are based on Savits (2006).

A.1 Efficient computation of derivatives of composite functions

The application of the SVD approach relies on Taylor expansion and this, in turns, creates the need for efficient computation of the derivatives of composite functions of the form $h(\varepsilon) = f(g(\varepsilon))$, where $f : \mathbb{R} \rightarrow \mathbb{R}$, $g : \mathbb{R}^N \rightarrow \mathbb{R}$, and ε denotes the N -dimensional vector of fundamental shocks. The generic M -th order Taylor expansion of h is

$$h(\varepsilon) \approx \sum_{\{q:|q|\leq M\}} \frac{1}{q!} h_q(0_N) \prod_{k=1}^N \varepsilon^{q_k}, \quad (\text{A1})$$

where $q = (q_1, \dots, q_N)$ is a vector of non-negative integers, $|q| = \sum_{n=1}^N q_n$, $q! = \prod_{n=1}^N (q_n!)$ and $h_q(0_N)$ denotes the derivative of h evaluated at 0_N , i.e.,

$$h_q(0_N) = \frac{\partial^{q_1+\dots+q_N} h}{\partial \varepsilon_1^{q_1} \dots \partial \varepsilon_N^{q_N}}(0_N). \quad (\text{A2})$$

To compute such derivatives we rely on the multi-dimensional version of the Faà di Bruno (1855, 1857) formula as derived in Savits (2006). We first introduce some notation. Let \mathbb{N}_0 denote the set of non-negative integers. If $q = (q_1, \dots, q_N) \in \mathbb{N}_0^N$, $\ell = (\ell_1, \dots, \ell_N) \in \mathbb{N}_0^N$, then we write $\ell \leq q$ provided $\ell_n \leq q_n$ for $n = 1, \dots, N$ and

$$\binom{q}{\ell} = \frac{q!}{\ell!(q-\ell)!} \quad (\text{A3})$$

Let $g_q(\varepsilon)$ denote the q -partial derivative of g , and $f_n(x)$ the n -th derivative of f with respect to the one-dimensional variable x . According to the multivariate Faà di Bruno formula, the q -partial derivative $h_q(\varepsilon)$ of the composite function $h(\varepsilon) = f(g(\varepsilon))$ can be written as

$$h_q(\varepsilon) = \sum_{n=1}^{|q|} f_n(g(\varepsilon)) \alpha_{q,n}(\varepsilon), \quad (\text{A4})$$

where $\alpha_{q,n}$ are homogeneous polynomials of degree n in the partial derivatives g_ℓ , $\ell \leq q$. To compute the generic derivative of h it is, therefore, sufficient to determine the polynomials $\alpha_{q,n}$. These can be computed efficiently by relying on the recursive relationship proved in Theorem 3.1 of Savits (2006), which we reproduce here.

Theorem 1 (Savits (2006)) For $q \geq 0_N$, $1 \leq j \leq N$ and $1 \leq n \leq |q| + 1$, we have:

$$\alpha_{q+e_j, n}(\varepsilon) = \sum_{\{\ell \in \mathbb{N}_0^N : 0_N \leq \ell \leq q, |\ell| \geq n-1\}} \binom{m}{\ell} g_{q+e_j-\ell}(\varepsilon) \alpha_{\ell, n-1}(\varepsilon), \quad (\text{A5})$$

where e_j is the N -dimensional unit vector with j -th component equal to 1 and we set

$$\alpha_{\ell, 0}(\varepsilon) = \begin{cases} 1, & \text{if } \ell = 0_N \\ 0, & \text{if } \ell \neq 0_N. \end{cases} \quad (\text{A6})$$

If the set $\{\ell \in \mathbb{N}_0^N : 0_N \leq \ell \leq q, |\ell| \geq n-1\}$ is empty, the polynomial $\alpha_{q+e_j, n}(\varepsilon)$ vanishes.

A.2 Efficient computation of multivariate moments of a normal random variable

Let $\varepsilon = (\varepsilon_1, \dots, \varepsilon_d) \sim \mathcal{N}(0_d, \Sigma)$ be a zero-mean d -dimensional normal random variable with covariance matrix Σ , with (i, j) element equal to σ_{ij} . Let $\lambda_{(\nu_1, \dots, \nu_d)}$ be its (ν_1, \dots, ν_d) -moment, where ν_1, \dots, ν_d are non-negative integers, i.e., $\lambda_{(\nu_1, \dots, \nu_d)} \equiv E[\varepsilon_1^{\nu_1} \dots \varepsilon_d^{\nu_d}]$. Then, from Theorem 5.1 in Savits (2006), we have the following recursive characterization of the multivariate moments of ε :

Theorem 2 (Savits (2006)) Set $\lambda_{(0, \dots, 0)} = 1$. Then, for all $(\nu_1, \dots, \nu_d) \geq (0, \dots, 0)$ and $1 \leq j \leq d$

$$\lambda_{(\nu_1, \dots, \nu_d)+e_j} \equiv E[\varepsilon_1^{\nu_1} \dots \varepsilon_j^{\nu_j+1} \dots \varepsilon_d^{\nu_d}] = \sum_{k=1}^d \sigma_{jk} \nu_k \lambda_{(\nu_1, \dots, \nu_d)} - e_k \quad (\text{A7})$$

where e_j is the d -dimensional unit vector with j -th component equal to 1 and all other components equal to zero.

B Proofs

Proof of Lemma 1

For the CARA utility function $U(W) = -\exp(-\alpha W)$ we have

$$U^{(m)}(W) = (-1)^{m+1} \alpha^m e^{-\alpha W}. \quad (\text{B8})$$

Therefore, for $\widetilde{W}_2 = \widetilde{W}_1 R_f + \frac{1}{\alpha} \mu' \Sigma^{-1} R_2$, we have

$$U^{(m)}(\widetilde{W}_2) = (-1)^{m+1} \alpha^m e^{-\alpha \widetilde{W}_1 R_f} e^{-Y_2}, \quad (\text{B9})$$

where $Y_2 = \mu' \Sigma^{-1} R_2 \sim \mathcal{N}(S^2, S^2)$ and $S^2 = \mu' \Sigma^{-1} \mu$ is the squared Sharpe ratio. Hence, using the binomial formula, we obtain

$$\begin{aligned} & E_1 \left[U^{(m)}(\widetilde{W}_2) \left(R_f + \frac{1}{\alpha \widetilde{W}_1} \mu' \Sigma^{-1} R_2 \right)^m \right] \\ &= (-1)^{m+1} \alpha^m e^{-\alpha \widetilde{W}_1 R_f} E_1 \left[e^{-Y_2} \left(R_f + \frac{1}{\alpha \widetilde{W}_1} Y_2 \right)^m \right] \\ &= (-1)^{m+1} \alpha^m e^{-\alpha \widetilde{W}_1 R_f} \sum_{k=0}^m \binom{m}{k} \left(\frac{1}{\alpha \widetilde{W}_1} \right)^k R_f^{m-k} E_1 \left[Y_2^k e^{-Y_2} \right]. \end{aligned} \quad (\text{B10})$$

Using the fact that $Y_2 \sim \mathcal{N}(S^2, S^2)$, we obtain

$$\begin{aligned} E_1 \left[Y_2^k e^{-Y_2} \right] &= \int \frac{1}{\sqrt{2\pi S^2}} y^k e^{-y} e^{-\frac{(y-S^2)^2}{2S^2}} dy \\ &= e^{-\frac{S^2}{2}} \int \frac{1}{\sqrt{2\pi S^2}} y^k e^{-\frac{y^2}{2S^2}} dy \\ &= e^{-\frac{S^2}{2}} S^k \int \frac{1}{\sqrt{2\pi}} z^k e^{-\frac{z^2}{2}} dz \\ &= e^{-\frac{S^2}{2}} S^k \phi_k, \end{aligned} \quad (\text{B11})$$

where ϕ_k is the k -th central moment of the standard normal distribution. Substituting $E_1 \left[Y_2^k e^{-Y_2} \right] = e^{-\frac{S^2}{2}} S^k \phi_k$ in (B10) yields the result and completes the proof. ■

C Tables

Table 1: Optimal portfolios under CRRA preferences

The table reports solutions to the CRRA portfolio choice problem (28) with CRRA utility and three risky assets. In the table, ω_i , $i = 1, 2, 3$ are the optimal portfolio weights in percent, γ is the constant of relative risk aversion, and M is the order of Taylor approximation used by the SVD method. The three risky assets we consider are the MSCI gross indices for USA, Europe, and Pacific obtained from MSCI-Barra. Under the assumption of normally distributed log excess returns, we estimate parameters using a time series from December 1969 to July 2006 as given in (33). The annual risk-free rate is assumed to be 5%. We report three sets of results. The first column reports the solution obtained by using quadrature to approximate expectations in the optimization problem. Moreover, we report results for two versions of the SVD method: the first uses a decomposition of excess returns while the second uses a decomposition of log excess returns. CEL is the certainty equivalent loss in annualized basis points.

		Quadrature	State Variable Decomposition					
			Excess Return			Log Excess Return		
			$M = 4$	$M = 5$	$M = 6$	$M = 4$	$M = 5$	$M = 6$
$\gamma = 5$	ω_1	23.91	26.40	24.84	24.14	23.88	23.88	23.91
	ω_2	22.82	24.27	23.61	22.93	22.81	22.81	22.82
	ω_3	10.70	12.65	11.32	10.86	10.69	10.69	10.70
	CEL		2.59	0.39	0.02	0.00	0.00	0.00
$\gamma = 10$	ω_1	11.94	12.74	12.20	11.99	11.93	11.93	11.94
	ω_2	11.34	11.79	11.57	11.36	11.35	11.35	11.34
	ω_3	5.30	5.94	5.48	5.34	5.30	5.30	5.30
	CEL		0.51	0.06	0.00	0.00	0.00	0.00
$\gamma = 15$	ω_1	7.95	29.56	8.09	7.97	7.94	7.94	7.95
	ω_2	7.54	20.15	7.66	7.55	7.55	7.55	7.54
	ω_3	3.52	50.29	3.62	3.54	3.52	3.52	3.52
	CEL		1743.40	0.03	0.00	0.00	0.00	0.00

Table 2: Comparison of BGSS and SVD methods

The table reports the certainty equivalent loss (CEL) associated with the BGSS and SVD solutions to the CRRA portfolio choice problem (28) with two risky assets. The benchmark used to compute the CEL is the solution obtained by computing expectations through Gauss-Hermite quadrature with 10 nodes in each dimension. The log excess returns on the two risky assets are assumed to be jointly normally distributed with means equal to 8% and 11% and standard deviations equal to 13% and 20%, respectively. The correlation between the log excess returns on the two assets takes 5 values: -0.5 , -0.25 , 0 , 0.25 , and 0.5 . The annualized risk-free rate is set equal to 5%, the coefficient of relative risk aversion is set equal to $\gamma = 10$, and M is the order of Taylor expansion ranging from 4 to 8. We impose no short sales and borrowing constraints. We report results for the BGSS method and two variants of the SVD method: one that uses a decomposition of excess returns and another that uses a decomposition of log excess returns. CEL is stated in annualized basis points.

Correlation	-0.5	-0.25	0	0.25	0.5
BGSS CEL					
$M = 4$	151.59	124.80	68.53	39.78	24.87
$M = 5$	1505.24	1332.87	1181.77	1086.91	1023.68
$M = 6$	81.99	81.86	40.25	20.29	10.99
$M = 7$	1505.24	1332.87	1181.77	1086.91	1023.68
$M = 8$	39.71	57.87	25.82	11.29	5.24
SVD (excess return decomposition) CEL					
$M = 4$	1.61	3.87	24.09	96.53	194.27
$M = 5$	0.45	1.36	19.29	47.79	9.41
$M = 6$	0.07	0.25	13.23	1.86	0.39
$M = 7$	0.02	0.11	2.95	0.61	0.19
$M = 8$	0.00	0.02	0.36	0.04	0.00
SVD (log excess return decomposition) CEL					
$M = 4$	0.00	0.00	0.77	0.10	0.01
$M = 5$	0.00	0.00	0.77	0.10	0.01
$M = 6$	0.00	0.00	0.00	0.00	0.00
$M = 7$	0.00	0.00	0.00	0.00	0.00
$M = 8$	0.00	0.00	0.00	0.00	0.00

Table 3: Certainty equivalent under CARA preferences and normal IID returns

The table reports the certainty equivalent returns obtained by the exact and the SVD approximate solutions for the dynamic portfolio choice problem with CARA preferences studied in subsection 4.1. There are three risky assets with i.i.d. normally distributed excess returns. The mean vector and the covariance matrix of the risky asset excess returns are provided in (46), while the annualized risk-free rate is set equal to 5%. The label Exact refers to the exact closed-form solution obtained in Proposition 1. The label SVD refers to the generic SVD approach with analytical computation of conditional expectations, as outlined in subsection 4.1 in which we approximate the value function V_t by a second-order polynomial and use a Taylor expansion of order $M = 4$.

		$W_0 = 1$	$W_0 = 1.25$	$W_0 = 1.5$	$W_0 = 1.75$	$W_0 = 2$
		Exact				
$\alpha = 2$	$T = 10$	10.066	9.081	8.417	7.939	7.578
	$T = 20$	9.284	8.480	7.930	7.531	7.227
	$T = 30$	8.670	7.998	7.535	7.196	6.937
$\alpha = 4$	$T = 10$	7.578	7.070	6.729	6.484	6.301
	$T = 20$	7.227	6.797	6.506	6.296	6.137
	$T = 30$	6.937	6.568	6.317	6.135	5.997
$\alpha = 6$	$T = 10$	6.729	6.386	6.157	5.993	5.870
	$T = 20$	6.506	6.211	6.013	5.871	5.763
	$T = 30$	6.317	6.062	5.889	5.765	5.672
		SVD-A				
$\alpha = 2$	$T = 10$	10.069	9.083	8.419	7.940	7.579
	$T = 20$	9.276	8.473	7.925	7.526	7.223
	$T = 30$	8.671	7.999	7.536	7.197	6.938
$\alpha = 4$	$T = 10$	7.577	7.069	6.729	6.484	6.300
	$T = 20$	7.236	6.804	6.512	6.301	6.142
	$T = 30$	6.934	6.565	6.315	6.133	5.996
$\alpha = 6$	$T = 10$	6.729	6.386	6.157	5.993	5.870
	$T = 20$	6.512	6.216	6.017	5.874	5.767
	$T = 30$	6.315	6.060	5.888	5.764	5.671

Table 4: Comparison of the BGSS, M1, M2, and SVD methods

The table reports the certainty equivalent loss in annual basis points for the four methods considered in section 4.2. The CEL of a method $i = \text{BGSS, M1, M2, SVD}$, is the difference between the certainty equivalent return associated with method i , computed according to (81), and the certainty equivalent return associated with exact closed-form solution given by (82). The initial wealth level W_0 takes values 0.5, 1, and 1.5 and risk aversion is set to $\alpha = 4$. The return on the risk free asset is 5% per annum, the annual expected returns of the risky assets, in excess of the risk-free rate, are 10% and 15% and their annual volatilities are 15% and 25%, respectively. The correlation between the excess returns on the two assets takes 3 values: -0.25, 0, and 0.25. The order of expansion in the Taylor approximation, M , takes values 4,6,8, and 10.

ρ		$W_0 = 0.5$			$W_0 = 1.0$			$W_0 = 1.5$		
		-0.25	0	0.25	-0.25	0	0.25	-0.25	0	0.25
BGSS	$M = 4$	239.48	114.42	59.84	78.74	33.47	15.84	46.33	19.11	8.81
	$M = 6$	157.05	61.02	26.45	34.45	9.76	3.19	16.57	4.18	1.22
	$M = 8$	107.67	35.39	14.36	13.86	2.63	0.76	4.68	0.67	0.16
	$M = 10$	77.12	23.77	10.61	5.44	1.04	0.42	1.12	0.17	0.06
M1	$M = 4$	2.57	4.44	3.83	1.69	0.07	0.00	3.41	0.50	0.12
	$M = 6$	30.59	16.92	10.13	2.55	1.26	0.69	0.35	0.19	0.10
	$M = 8$	37.81	19.41	11.16	3.95	1.68	0.84	0.81	0.32	0.15
	$M = 10$	39.35	19.83	11.30	4.20	1.73	0.85	0.90	0.34	0.15
M2	$M = 4$	105.22	43.65	20.10	57.43	23.43	10.67	39.57	16.03	7.27
	$M = 6$	28.78	6.59	1.70	15.72	3.54	0.90	10.84	2.42	0.62
	$M = 8$	4.38	0.38	0.04	2.39	0.20	0.02	1.65	0.14	0.02
	$M = 10$	0.26	0.01	0.00	0.14	0.00	0.00	0.10	0.00	0.00
SVD	$M = 4$	21.96	4.41	1.43	12.00	2.37	0.76	8.27	1.62	0.52
	$M = 6$	0.29	0.05	0.01	0.16	0.03	0.01	0.11	0.02	0.00
	$M = 8$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	$M = 10$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 5: Strategic asset allocation with predictable returns: SVD versus Quadrature

The table reports the portfolio weights in the nominal T-bond (ω_b) and stock (ω_s), and the consumption-to-wealth ratio (c) for the strategic asset allocation problem discussed in Section 5. Panel A considers a 15-year horizon problem while Panel B presents a 30-year problem. Each sub-panel reports the portfolio allocation at different quartiles in the distribution of the variable mentioned in the heading of the panel, assuming that all the remaining 5 variables are kept at their median. The columns labeled “Q” contain the solution for the quadrature method while the columns labeled “SVD” contain the solution for the SVD method.

	Panel A: Investment horizon 15 years					Panel B: Investment horizon 30 years						
	Q	SVD	Q	SVD	Q	SVD	Q	SVD	Q	SVD	Q	SVD
	p_{25}		p_{50}		p_{75}		p_{25}		p_{50}		p_{75}	
Log real T-bill return (r_1)												
ω_b	21.38	21.31	48.03	48.29	51.69	51.94	32.01	31.80	46.48	46.92	50.25	50.70
ω_s	53.52	53.24	51.97	51.71	48.31	48.06	56.29	55.83	53.52	53.08	49.75	49.30
c	8.68	8.68	8.76	8.75	8.84	8.84	6.13	6.12	6.17	6.17	6.22	6.21
Log excess real T-bond return (r_2)												
ω_b	50.54	50.79	48.03	48.29	45.51	45.78	49.01	49.45	46.48	46.92	43.95	44.39
ω_s	49.46	49.21	51.97	51.71	54.49	54.22	50.99	50.55	53.52	53.08	56.05	55.61
c	8.74	8.73	8.76	8.75	8.78	8.78	6.15	6.15	6.17	6.17	6.19	6.19
Log excess real stock return (r_3)												
ω_b	11.29	11.16	48.03	48.29	49.90	50.15	21.99	21.78	46.48	46.92	48.36	48.81
ω_s	50.69	50.40	51.97	51.71	50.10	49.85	53.35	52.88	53.52	53.08	51.64	51.19
c	8.77	8.77	8.76	8.75	8.75	8.74	6.19	6.18	6.17	6.17	6.16	6.15
Short term nominal interest rate (z_1)												
ω_b	48.05	48.32	48.03	48.29	48.06	48.29	46.03	46.50	46.48	46.92	47.00	47.41
ω_s	51.95	51.68	51.97	51.71	51.94	51.71	53.97	53.50	53.52	53.08	53.00	52.59
c	8.52	8.51	8.76	8.75	9.01	9.01	5.95	5.95	6.17	6.17	6.42	6.41
Dividend yield (z_2)												
ω_b	66.33	66.56	48.03	48.29	29.74	30.02	65.43	65.85	46.48	46.92	27.61	28.11
ω_s	33.67	33.44	51.97	51.71	70.26	69.98	34.57	34.15	53.52	53.08	72.39	71.89
c	8.63	8.63	8.76	8.75	8.95	8.95	6.07	6.07	6.17	6.17	6.33	6.33
Yield spread (z_3)												
ω_b	0.00	0.00	48.03	48.29	54.83	55.04	0.00	0.00	46.48	46.92	53.54	53.91
ω_s	50.19	49.90	51.97	51.71	45.17	44.96	52.99	52.51	53.52	53.08	46.46	46.09
c	8.72	8.71	8.76	8.75	8.83	8.83	6.13	6.12	6.17	6.17	6.24	6.23

Table 6: Strategic asset allocation with predictable returns: EIS comparative statics

The table reports the portfolio weights in the nominal T-bond (ω_b) and stock (ω_s), and the consumption-to-wealth ratio (c) for the strategic asset allocation problem discussed in Section 5. Panel A considers a 15-year horizon problem while Panel B presents a 30-year problem. Each panel considers two different values for the elasticity of intertemporal substitution (EIS). Each sub-panel reports the portfolio allocation at different quartiles in the distribution of the variable mentioned in the heading of the panel, assuming that all the remaining 5 variables are kept at their median.

		Panel A: Investment horizon 15 years					Panel B: Investment horizon 30 years						
EIS		0.75	1.25	0.75	1.25	0.75	1.25	0.75	1.25	0.75	1.25		
		p_{25}					p_{75}						
		p_{50}					p_{25}						
		p_{75}					p_{50}						
		p_{25}					p_{75}						
ω_b		20.84	20.01	48.41	48.60	52.08	52.27	31.02	29.14	46.88	47.25	50.66	51.03
ω_s		53.07	52.80	51.59	51.40	47.92	47.73	55.78	55.22	53.12	52.75	49.34	48.97
c		9.11	10.00	9.15	9.95	9.19	9.91	6.57	7.51	6.59	7.48	6.62	7.45
Log real T-bill return (r_1)													
ω_b		50.93	51.13	48.41	48.60	45.90	46.07	49.41	49.79	46.88	47.25	44.36	44.71
ω_s		49.07	48.87	51.59	51.40	54.10	53.93	50.59	50.21	53.12	52.75	55.64	55.29
c		9.14	9.97	9.15	9.95	9.16	9.94	6.58	7.49	6.59	7.48	6.61	7.47
Log excess real T-bond return (r_2)													
ω_b		10.71	9.82	48.41	48.60	50.29	50.49	20.95	18.99	46.88	47.25	48.76	49.16
ω_s		50.23	49.95	51.59	51.40	49.71	49.51	52.84	52.28	53.12	52.75	51.24	50.84
c		9.16	9.95	9.15	9.95	9.14	9.96	6.60	7.47	6.59	7.48	6.59	7.49
Log excess real stock return (r_3)													
ω_b		48.52	48.80	48.41	48.60	48.38	48.46	46.54	47.05	46.88	47.25	47.30	47.54
ω_s		51.48	51.20	51.59	51.40	51.62	51.54	53.46	52.95	53.12	52.75	52.70	52.46
c		9.02	10.09	9.15	9.95	9.28	9.82	6.48	7.61	6.59	7.48	6.72	7.35
Short term nominal interest rate (z_1)													
ω_b		66.70	66.83	48.41	48.60	30.09	30.32	65.82	66.06	46.88	47.25	28.05	28.52
ω_s		33.30	33.17	51.59	51.40	69.91	69.68	34.18	33.94	53.12	52.75	71.95	71.48
c		9.08	10.02	9.15	9.95	9.25	9.85	6.54	7.54	6.59	7.48	6.68	7.39
Dividend yield (z_2)													
ω_b		0.00	0.00	48.41	48.60	55.17	55.35	0.00	0.00	46.88	47.25	53.85	54.20
ω_s		49.70	49.40	51.59	51.40	44.83	44.65	52.41	51.84	53.12	52.75	46.15	45.80
c		9.13	9.98	9.15	9.95	9.19	9.91	6.57	7.51	6.59	7.48	6.63	7.44
Yield spread (z_3)													

Table 7: Strategic asset allocation with predictable returns: Hedging demands

The table reports the portfolio weights in the nominal T-bill (ω_f), nominal T-bond (ω_b) and stock (ω_s) for the strategic asset allocation problem discussed in Section 5. Panel A considers a 15-year horizon and Panel B is the solution to a 30-year problem. Each sub-panel reports the portfolio allocation at different quartiles in the distribution of the variable mentioned in the heading of the panel, assuming that all the remaining 5 variables are kept at their median. The columns labeled “M” represents the myopic demand while the columns labeled “H” are the hedging demands.

	Panel A: Investment horizon 15 years					Panel B: Investment horizon 30 years					
	M	H	M	H	M	H	M	H	M	H	
	p_{25}		p_{50}		p_{75}		p_{25}		p_{50}		p_{75}
	Log real T-bill return (r_1)										
ω_f	63.46	-38.01	0.00	0.00	0.00	63.46	-51.09	0.00	0.00	0.00	0.00
ω_b	0.00	21.31	60.56	-12.28	63.51	-11.57	0.00	31.80	60.56	-13.65	63.51
ω_s	36.54	16.70	39.44	12.28	36.49	11.57	36.54	19.29	39.44	13.65	36.49
	Log excess real T-bond return (r_2)										
ω_f	0.00	-0.00	0.00	0.00	37.13	-37.13	0.00	-0.00	0.00	0.00	37.13
ω_b	62.97	-12.18	60.56	-12.28	24.08	21.69	62.97	-13.53	60.56	-13.65	24.08
ω_s	37.03	12.18	39.44	12.28	38.78	15.44	37.03	13.53	39.44	13.65	38.78
	Log excess real stock return (r_3)										
ω_f	65.53	-27.09	0.00	-0.00	0.00	0.00	65.53	-40.20	0.00	-0.00	0.00
ω_b	0.00	11.16	60.56	-12.28	62.40	-12.25	0.00	21.78	60.56	-13.65	62.40
ω_s	34.47	15.94	39.44	12.28	37.60	12.25	34.47	18.41	39.44	13.65	37.60
	Short term nominal interest rate (z_1)										
ω_f	0.00	0.00	0.00	-0.00	0.00	0.00	0.00	0.00	0.00	-0.00	0.00
ω_b	61.39	-13.08	60.56	-12.28	59.73	-11.44	61.39	-14.89	60.56	-13.65	59.73
ω_s	38.61	13.08	39.44	12.28	40.27	11.44	38.61	14.89	39.44	13.65	40.27
	Dividend yield (z_2)										
ω_f	21.19	-21.19	0.00	0.00	0.00	0.00	21.19	-21.19	0.00	0.00	0.00
ω_b	56.68	9.88	60.56	-12.28	44.96	-14.94	56.68	9.17	60.56	-13.65	44.96
ω_s	22.13	11.31	39.44	12.28	55.04	14.94	22.13	12.02	39.44	13.65	55.04
	Yield spread (z_3)										
ω_f	66.57	-16.47	0.00	-0.00	0.00	0.00	66.57	-19.09	0.00	0.00	0.00
ω_b	0.00	0.00	60.56	-12.28	66.52	-11.48	0.00	0.00	60.56	-13.65	66.52
ω_s	33.43	16.47	39.44	12.28	33.48	11.48	33.43	19.09	39.44	13.65	33.48

References

- Aït-Sahalia, Y., J. Cacho-Diaz, and T. R. Hurd, 2008, "Portfolio Choice with Jumps: A Closed-Form Solution," *Annals of Applied Probability*, forthcoming.
- Apostol, T. M., 1974, *Mathematical Analysis*, Addison-Wesley.
- Balduzzi, P., and A. Lynch, 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.
- Bhamra, H., and R. Uppal, 2006, "The Role of Risk Aversion and Intertemporal Substitution in Dynamic Consumption-Portfolio Choice with Recursive Utility," *Journal of Economic Dynamics and Control*, 30, 967–991.
- Brandt, M. W., 1999, "Estimating Portfolio and Consumption Choice: A Conditional Euler Equation Approach," *Journal of Finance*, 54, 1609–1646.
- Brandt, M. W., A. Goyal, P. Santa-Clara, and J. R. Stroud, 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., E. S. Schwartz, and R. Lagnado, 1997, "Strategic Asset Allocation," *Journal of Economic Dynamics and Control*, 21, 1377–1403.
- Campbell, J. Y., Y. L. Chan, and L. M. Viceira, 2003, "A Multivariate Model of Strategic Asset Allocation," *Journal of Financial Economics*, 67, 41–80.
- Campbell, J. Y., J. Cocco, F. Gomes, P. J. Maenhout, and L. M. Viceira, 2001, "Stock Market Mean Reversion and the Optimal Equity Allocation of a Long-Lived Investor," *European Finance Review*, 5, 269–292.
- Campbell, J. Y., and L. M. Viceira, 1999, "Consumption and Portfolio Decisions when Expected Returns are Time Varying," *Quarterly Journal of Economics*, 114, 433–495.
- Chacko, G., and L. M. Viceira, 2005, "Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets," *Review of Financial Studies*, 18, 1369–1402.
- Cvitanović, J., A. Lazrak, L. Martellini, and F. Zapatero, 2006, "Dynamic Portfolio Choice with Uncertainty and the Economic Value of Analysts Recommendations," *Journal of Economic Theory*, 19, 1113–1156.
- Das, S. R., and R. K. Sundaram, 2002, "A Numerical Algorithm for Consumption-Investment Problems," *International Journal of Intelligent Systems in Accounting, Finance and Management*, 11, 55–69.

- Detemple, J. B., R. Garcia, and M. Rindisbacher, 2003, “A Monte Carlo Method for Optimal Portfolios,” *Journal of Finance*, 58, 410–446.
- Epstein, L. G., and S. Zin, 1989, “Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–969.
- Epstein, L. G., and S. Zin, 1991, “Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *Journal of Political Economy*, 99, 263–286.
- Faà di Bruno, F., 1855, “Sullo Sviluppo delle Funzioni,” *Annali di Scienze Matematiche e Fisiche*, 6, 479–480.
- Faà di Bruno, F., 1857, “Note Sur une Nouvelle Formule de Calcul Différentiel,” *Quarterly Journal of Pure and Applied Mathematics*, 1, 359–360.
- Garlappi, L., and G. Skoulakis, 2008a, “Numerical Solutions to Dynamic Portfolio Problems: The Case for Value Function Iteration using Taylor Expansion,” *Computational Economics*, forthcoming.
- Garlappi, L., and G. Skoulakis, 2008b, “Taylor Series Approximations to Expected Utility and Portfolio Choice,” Working Paper, University of Texas at Austin and University of Maryland.
- Glasserman, P., 2004, *Monte Carlo Methods in Financial Engineering*, Springer-Verlag.
- Hakansson, N. H., 1971, “Capital Growth and the Mean-Variance Approach to Portfolio Selection,” *Journal of Financial and Quantitative Analysis*, 6, 517–557.
- Haykin, S., 1999, *Neural Networks: A Comprehensive Foundation*, Prentice Hall, 2nd edn.
- Hlawitschka, W., 1994, “The Empirical Nature of Taylor-Series Approximations to Expected Utility,” *The American Economic Review*, 84, 713–71.
- Judd, K., 1998, *Numerical Methods in Economics*, MIT Press.
- Kim, T. S., and E. Omberg, 1996, “Dynamic Nonmyopic Portfolio Behavior,” *Review of Financial Studies*, 9, 141–61.
- Kogan, L., and R. Uppal, 2003, “Risk Aversion and Optimal Portfolio Policies in Partial and General Equilibrium Economies,” CEPR Discussion Paper 3306 and NBER Working paper 8609.
- Kroll, Y., H. Levy, and H. M. Markowitz, 1984, “Mean-Variance Versus Direct Utility Maximization,” *Journal of Finance*, 39, 47–75.
- Liu, J., 2007, “Portfolio Selection in Stochastic Environments,” *Review of Financial Studies*, 20.
- Loistl, O., 1976, “The Erroneous Approximation of Expected Utility by Means of a Taylor’s Series Expansion: Analytic and Computational Results,” *The American Economic Review*, 66, 904–910.
- Longstaff, F. A., and E. S. Schwartz, 2001, “Valuing American Options by Simulation: A Simple Least-Squares Approach,” *Review of Financial Studies*, 14, 113–147.
- Markowitz, H. M., 1991, “Foundations of Portfolio Theory,” *Journal of Finance*, 46, 469–477.

- Merton, R. C., 1969, "Lifetime Portfolio Choice: The Continuous-Time Case," *Review of Economics and Statistics*, 51, 247–257.
- Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.
- Merton, R. C., 1973, "An Intertemporal Asset Pricing Model," *Econometrica*, 41, 867–888.
- Niederreiter, H., 1992, *Random Number Generation and Quasi-Monte Carlo Methods*, Society for Industrial and Applied Mathematics.
- Pulley, L. B., 1981, "A General Mean-Variance Approximation to Expected Utility for Short Holding Periods," *Journal of Financial and Quantitative Analysis*, 16, 361–373.
- Pulley, L. B., 1983, "Mean-Variance Approximations to Expected Logarithmic Utility," *Operations Research*, 31, 685–696.
- Samuelson, P., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 239–246.
- Samuelson, P. A., 1970, "The Fundamental Approximation Theorem of Portfolio Analysis in Terms of Means, Variances and Moments," *Review of Economic Studies*, 37, 537–542.
- Savits, T. H., 2006, "Some Statistical Applications of Faà di Bruno," *Journal of Multivariate Analysis*, 97, 2131–2140.
- Skoulakis, G., 2008, "A Recursive Formula for Computing Central Moments of a Multivariate Lognormal Distribution," *The American Statistician*, 62, 147–150.
- Tsitsiklis, J. N., and B. Van Roy, 2001, "Regression Methods for Pricing Complex American-Style Options," *IEEE Transactions on Neural Networks*, 12, 694–703.
- Vissing-Jørgensen, A., 2002, "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110, 825–83.
- Wachter, J. A., 2002, "Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets," *Journal of Financial and Quantitative Analysis*, 37, 63–91.