

On Approximating the Gradient of the Value Function

Alexandre Dmitriev^{1,2}

¹Department of Economics, University of Auckland, Auckland, New Zealand

²Centre for Applied Macroeconomic Analysis (CAMA), Australian National University, Canberra, Australia

Email: a.dmitriev@auckland.ac.nz

How to cite this paper: Dmitriev, A. (2019) On Approximating the Gradient of the Value Function. *Theoretical Economics Letters*, 9, 126-138.

<https://doi.org/10.4236/tel.2019.91011>

Received: December 27, 2018

Accepted: January 29, 2019

Published: February 1, 2019

Copyright © 2019 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

The optimality conditions for macroeconomic problems with limited commitment often contain partial derivatives of the optimal value function, corresponding to the outside option. This paper contributes to the literature on recursive contracts by proposing an algorithm for approximating the gradient of the value function using simulation-based methods. Our method combines numerical solution and simulation of the model, Monte-Carlo integration and numerical differentiation. It does not suffer from the curse of dimensionality and is therefore convenient for models involving many state variables. The algorithm inherits the speed and accuracy limitations of the numerical solution method it relies on. Our accuracy analysis is limited to a few classical examples from macroeconomic literature.

Keywords

Dynamic Participation Constrains, Numerical Algorithm, Recursive Contracts

1. Introduction

The purpose of this paper is to propose a simple algorithm for computing partial derivatives of the optimal value function. Macroeconomic problems involving incentive compatibility constraints have received wide attention in the literature due to recent advances in dynamic optimization techniques (see [1] [2] and references therein). Often the optimality conditions for this class of problems involve partial derivatives with respect to endogenous state variables of the optimal value function corresponding to the dynamic programming formulation of an outside option. Although many numerical methods can provide an approximation for the value function, there is no reason to believe that a derivative of this

approximation will be close in any sense to the actual value of the derivative. In this note we suggest an algorithm for accurately computing these partial derivatives by simulation.

This issue has been previously considered in [3], in the context of a stochastic growth model with capital accumulation under one-sided lack of commitment. To circumvent the problem of finding the values of the derivatives in [3], the authors proposed a method based on the ideas of Benveniste and Scheinkman [4]. Unfortunately, their method has limited applicability since it depends on the availability of an analytical solution for the derivatives as conditional expectations of the known functions of the model solution. This paper proposes a simple algorithm to fill this gap in the literature.

In order to be able to use finite differences to approximate the gradient at a given point, one would need to know the values of the optimal value function at a certain set of points. Our algorithm obtains approximations of these values with arbitrary precision. Moreover, achieving this accuracy is feasible for all points in the state-space which have economic relevance.

The initial step of our algorithm involves obtaining numerical solution to a problem using a procedure which satisfies three criteria. First, it approximates some unknown function with flexible functional forms of finite elements. Second, it can deliver an accurate solution as the number of the finite elements in the function goes to infinity. Third and last, the resulting numerical solution must be such that it can be formulated as a set of policy functions approximated with flexible functional forms. The next step involves using Monte-Carlo integration in order to evaluate the conditional expectation of the discounted sum of future instantaneous utilities. The final step involves applying the method of finite differences to approximate the values of the partial derivatives of the value function.

The attractive features of the algorithm include its rather wide scope of applicability and simplicity of implementation. It can be used to study the questions of risk sharing under imperfect enforcement of contracts, as well as partnerships with limited commitment when several state variables appear in the model corresponding to the outside option. Such models may include habit formation preferences, several types of capital, or reputational co-state variables. The suggested method is computationally inexpensive. It does not suffer from the curse of dimensionality and therefore it is particularly convenient for models involving many state variables.

The rest of the paper is organized as follows. Section 2 discusses an example, where our algorithm proves to be useful. Section 3 sketches the idea behind the algorithm. Section 4 deals with implementation of the algorithm, while Section 5 compares it with some available alternatives. Section 6 concludes.

2. Applicability of the Algorithm: An Example

To fix ideas, we start with an example of a macroeconomic model where our computational algorithm proves to be useful. The key feature of this example is

that solving it boils down to designing an optimal social contract which takes into account not only technological but also incentive and legal constraints. Our example illustrates the need for computing the gradient of the value function and its practical implementation. Furthermore, it shows that our algorithm is applicable to some widely used models, to which the method in [3] cannot be applied.

Consider a model of international risk sharing, which distinguishes itself from the canonical model [5] in two respects. First, as in [6], we introduce a friction in the credit markets. We assume that the international loans are feasible only to the extent to which they can be enforced by the threat of exclusion from participation in any other international risk sharing arrangement. Second, we incorporate habit formation preferences into the model. The motivation for doing this is threefold. Habits help us illustrate the features of the algorithm by expanding the set of endogenous states in the model. Habit formation preferences tend to improve performance of the international business cycle models [7] [8]. Finally, empirical studies suggest that habit formation is consistent with the observed consumption behavior [9]. To simplify the exposition, we assume inelastic labor supply.

The planner's problem is to choose the sequences of consumption $\{c_{it}\}$ and investment $\{i_{it}\}$ to maximize a weighted sum of utilities

$$\max_{\{c_{it}, i_{it}\}} E_0 \sum_{i=1}^I \lambda_i \sum_{t=0}^{\infty} \beta^t u(c_{it}, h_{it}) \quad (1)$$

subject to an aggregate feasibility constraint

$$\sum_{i=1}^I c_{it} + \sum_{i=1}^I i_{it} = \sum_{i=1}^I f(k_{it}, \theta_{it}), \quad (2)$$

individual participation constraints for each $i = 1, \dots, I$,

$$E_t \sum_{j=0}^{\infty} \beta^j u(c_{it+j}, h_{it+j}) \geq V_i^a(k_{it}, h_{it}, \theta_{it}), \quad (3)$$

the equations of motion for the capital,

$$k_{it+1} = (1 - \delta)k_{it} + i_{it}, \quad (4)$$

the laws of motion for habits,

$$h_{it+1} = h_{it} + \lambda(c_{it} - h_{it}), \quad (5)$$

and non-negativity constraints $c_{it}, i_{it} \geq 0$. We assume that productivity shocks θ_t follow a first order stationary vector autoregressive process, and that the initial values for the state variables $k_{i0}, h_{i0}, \theta_{i0}$, and the initial non-negative weights, λ_i , are given. In addition, the usual restrictions apply to the discount factor, $\beta \in (0, 1)$, the capital depreciation rate, $\delta \in (0, 1)$, and the persistence of habits, $\lambda \in (0, 1)$.

The outside option $V_i^a(k_{it}, h_{it}, \theta_{it})$ in the participation constraint (3) represents the optimal value function corresponding to the autarkic environment.

$$\max_{\{c_{it}, i_{it}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_{it}, h_{it}) \tag{6}$$

subject to

$$\begin{aligned} c_{it} + i_{it} &= f(k_{it}, \theta_{it}), \\ k_{it+1} &= f(k_{it}, \theta_{it}) - c_{it} + (1 - \delta)k_{it}, \\ h_{it+1} &= h_{it} + \lambda(c_{it} - h_{it}), \end{aligned}$$

with the initial values being equal to the values of the state variables $k_{it}, h_{it}, \theta_{it}$ at the moment of deviation from the optimal plan.

In addition to Equations (2)-(5), the optimal allocations, for all $i, s = 1, \dots, I$, must satisfy the risk sharing condition,

$$\frac{\Lambda_{i,t}}{\Lambda_{s,t}} = \frac{\xi_{st}}{\xi_{it}},$$

where

$$\begin{aligned} \Lambda_{i,t} = & u_c(i, t) + \lambda \beta E_t \sum_{j=0}^{\infty} \beta^j (1 - \lambda)^j \left[\frac{\xi_{it+1+j}}{\xi_{it}} u_h(i, t+1+j) \right. \\ & \left. - \frac{\mu_{it+1+j}}{\xi_{it}} \frac{\partial V_i^a}{\partial h_{it+1+j}}(i, t+1+j) \right], \end{aligned} \tag{7}$$

the intertemporal condition,

$$\Lambda_{i,t} = \beta E_t \left[\frac{\xi_{it+1+j}}{\xi_{it}} \Lambda_{i,t+1} (f_k(k_{it+1}, \theta_{it+1}) + 1 - \delta) - \frac{\mu_{it+1}}{\xi_{it}} \frac{\partial V_i^a}{\partial k_{it+1}}(i, t+1) \right], \tag{8}$$

the complementary slackness condition,

$$\mu_{it} \left[E_t \sum_{j=0}^{\infty} \beta^j u(c_{it+j}, h_{it+j}) - V_i^a(k_{it}, h_{it}, \theta_{it}) \right] = 0,$$

and the law of motion for the co-state variables $M_{it+1} = M_{it} + \mu_{it}$, where $\xi_{it} = \lambda_i + M_{it+1}$, $\mu_{it} \geq 0$, and $M_{i0} = 0$. In the equations above $u_c(i, t)$ denotes $\frac{\partial u(c_{it}, h_{it})}{\partial c_{it}}$, and similar abbreviations apply to other terms¹.

The gradient of the optimal value function V_i^a enters the intertemporal condition (8) and the risk-sharing condition (7). Approximation of this gradient is the purpose of the algorithm proposed in [3] and in this paper. Because both algorithms can approximate $\frac{\partial V_i^a}{\partial k_{it}}(k_{it}, h_{it}, \theta_{it})$, we will use that fact to compare their accuracy in Section 1. Our approach can also approximate $\frac{\partial V_i^a}{\partial h_{it}}(k_{it}, h_{it}, \theta_{it})$,

whereas [3] cannot, because the analytical solution for this partial derivative as an expectation of the known functions of the model's solution is not available. This will be further discussed in Section 1.

¹Derivation of the optimality condition is available upon request.

3. The Algorithm

Typically, in the models with participation constraints the reservation value is the value function of the outside alternative, evaluated at the current values of the endogenous state variables, \bar{x} , and exogenous shocks, \bar{s} . Consider the optimal value function at a point (\bar{x}, \bar{s}) as an outcome of a standard optimization problem for the outside alternative:

$$V(\bar{x}, \bar{s}) = \max_{\{a_t\}} E_0 \sum_{t=0}^{\infty} \beta^t r(x_t, a_t, s_t) \quad (9)$$

subject to

$$\begin{aligned} x_{t+1} &= l(x_t, a_t, s_t), \quad a_t \in A(x_t, s_t), \\ x_0 &= \bar{x}, s_0 = \bar{s}, \end{aligned} \quad (10)$$

where r is an instantaneous utility function, $\beta \in (0, 1)$ the discount factor, $\{s_t\}$ an exogenous Markov stochastic process, x_t a vector of endogenous state variables, a_t a vector of control variables, A a feasibility correspondence and l the law of motion for the endogenous state variables. The functional equation to this problem can be derived using the standard dynamic programming techniques. It yields a time invariant policy function f such that the optimal allocations satisfy $a_t = f(x_t, s_t)$.

The purpose of our algorithm is to find a pointwise approximation to the partial derivative $\frac{\partial}{\partial x_i} V(\bar{x}, \bar{s})$ of the value function with respect to its i -th argument. The algorithm takes the following three steps:

Step I (Numerical Solution) Solve the model in (9) with a spectral method and formulate the solution in terms of approximated policy functions $a_t = \hat{f}(\omega; x_t, s_t)$, which depend on the state variables and some coefficients, ω .

Step II (Monte Carlo Integration) Simulate N sequences of the realizations of the stochastic process $\{s_t^n\}_{t=1}^T$ of size T with a starting value $s_0^n = \bar{s}$, for all $n = 1, \dots, N$. For a each sequence $\{s_t^n\}_{t=0}^T$, simulate the series of the endogenous variables $\{x_t^n, a_t^n\}_{t=0}^T$ using approximated policy functions \hat{f} , the equations for motion for the state variables (10), and the initial values $x_0^n = \bar{x}$. Using the simulated series calculate the discounted sums of the instantaneous returns and average over N :

$$V(\bar{x}, \bar{s}) \simeq \frac{1}{N} \sum_{n=1}^N \sum_{t=0}^T \beta^t r(x_t^n, a_t^n, s_t^n).$$

Step III (Numerical Differentiation) Repeat Step II to obtain approximations of the value function at two points, for instance $V(\bar{x} + \epsilon \iota_i, \bar{s})$ and $V(\bar{x} - \epsilon \iota_i, \bar{s})$, where ι_i denotes a conformable vector of zeros with one on its i -th coordinate, and ϵ is a small positive number. Calculate the value of the partial derivative using, for example, Stirling's finite difference formula:

$$\frac{\partial}{\partial x_i} V(\bar{x}, \bar{s}) \simeq \frac{V(\bar{x} + \epsilon \iota_i, \bar{s}) - V(\bar{x} - \epsilon \iota_i, \bar{s})}{2\epsilon}.$$

The optimal choice of the method for calculating the derivatives in Step III is problem specific and its accuracy depends on the smoothness of the value function. The approaches available include a variety of difference formulas, Richardson Extrapolation, or curve fitting with cubic splines. These are described at length in the standard numerical methods texts such as [10] [11] [12].

A brief note should be made at this point on the accuracy of the algorithm. In principle, arbitrary accuracy of the approximation can be achieved, by simultaneously increasing the dimension of the approximating family of functions in Step I, increasing the size of Monte Carlo iterations in Steps I and II, and decreasing the denominator ϵ in Step III. In practical applications, however, there are several sources of the approximation errors. First, in order to obtain the values of the optimal value function at a point, one relies on the approximations of the policy functions implied by the numerical solution to the model. Second, because we consider stochastic models, there is an additional error stemming from the evaluation of the integral in the computation of expected discounted returns. Finally, numerical differentiation introduces two more sources of error: the truncation error and the roundoff error. The truncation error comes from omitting higher order terms in the Taylor series expansion. The roundoff error is associated with storing real numbers in computer's floating-point format. Section 0 discusses some practical accuracy issues in the context of an example.

4. Implementation of the Algorithm

This section describes a practical computational strategy for implementing the algorithm using the example from Section 2. The optimality conditions include partial derivatives of the value function corresponding to the dynamic programming formulation of the agents outside option, *i.e.* autarky. The functional equation for the autarkic problem is:

$$V(k, h, \theta) = \max_{(c, i) \in A(k, \theta)} \left\{ u(c, h) + \beta E[V(k', h', \theta') | (k, h, \theta)] \right\}$$

$$h' = h + \lambda(c - h),$$

$$k' = (1 - \delta)k + i,$$

$$A(k, \theta) = \left\{ (c, i) \in \mathbb{R}_+^2 : c + i = f(k, \theta) \right\}.$$

The objective of the algorithm is to find the values of the partials $V_h(\cdot)$ and $V_k(\cdot)$ at a point $(\bar{k}, \bar{h}, \bar{\theta})$, which is likely to happen in equilibrium. Since the analytical expression for these derivatives is in general unavailable, we have no choice but to rely on numerical differentiation. Another complication which arises here is that the closed form solution to the optimal value function is generally unavailable too. Hence, one needs to approximate value function at two points, e.g. $V(\bar{k} + \varepsilon, \bar{h}, \bar{\theta})$ and $V(\bar{k} - \varepsilon, \bar{h}, \bar{\theta})$ with arbitrary accuracy in order to be able to use the finite differencing approach.

The first step of the algorithm involves solving the model with a spectral me-

thod which can approximate the policy functions with arbitrary accuracy. This example relies on a version of stochastic simulation algorithm, which formulates the solution in terms of approximated policy functions. The Euler equation for the problem is given by:

$$\Lambda_t = \beta E_t \left[\Lambda_{t+1} \left(f_k(k_{t+1}, \theta_{t+1}) + 1 - \delta \right) \right],$$

where marginal utility of consumption is

$$\Lambda_t = u_c(c_t, h_t) + \beta \lambda E_t \left[\sum_{j=0}^{\infty} \beta^j (1 - \lambda)^j u_h(c_{t+j+1}, h_{t+j+1}) \right].$$

To simplify the exposition we will consider the case of non-persistent habits, which corresponds to $\lambda = 1$. We assume the functional forms standard in the growth literature. The instantaneous utility function is given by

$$u(c_t, h_t) = \frac{(c_t - bh_t)^{1-\sigma}}{1-\sigma}, \text{ where } b \in (0,1) \text{ and } \sigma > 0. \text{ The production function}$$

is Cobb-Douglas and is given by $f(k_t, \theta_t) = \theta_t k_t^\alpha$. The stochastic process for productivity is $\log \theta_t = \rho \log \theta_{t-1} + \varepsilon_t$, where $\rho \in (0,1)$, and $\{\varepsilon_t\}$ are independent normally distributed random variables with zero mean and variance σ_ε^2 . In this example, we restrict attention to one particular set of the parameters which are summarized in **Table 1**.

The algorithm follows the three steps: 1) numerical solution; 2) Monte Carlo integration; 3) numerical differentiation.

Step I (Numerical Solution) The sequences of optimal allocations $\{c_t, h_{t+1}, k_{t+1}\}_{t=0}^{\infty}$ must satisfy the following system of stochastic difference equations:

$$(c_t - bh_t)^{-\sigma} = \beta E_t \left[b(c_{t+1} - bh_{t+1})^{-\sigma} \times \left(1 + (\alpha \theta_{t+1} k_{t+1}^{\alpha-1} + 1 - \delta) \left(\frac{1}{b} - \beta \left(\frac{c_{t+2} - bh_{t+2}}{c_{t+1} - bh_{t+1}} \right)^{-\sigma} \right) \right) \right], \tag{11}$$

$$k_{t+1} = \theta_t k_t^\alpha + (1 - \delta) k_t - c_t, \tag{12}$$

$$h_{t+1} = c_t. \tag{13}$$

For the expositional purpose, we solve the model with a version of a stochastic simulation algorithm, which is easiest to implement (see e.g. [13]). It takes the following steps:

1) Fix the initial conditions and draw a series of $\{\theta_t\}_{t=1}^T$ that obeys the law of motion for the exogenous shocks with T sufficiently large. To ensure sufficient accuracy of solution we chose $T = 50000$ for all the numerical examples considered. The computational burden of this is still rather low since the model needs to be solved only once.

2) Substitute the conditional expectations in (11) with the flexible functional forms that depend on the state variables k_t, h_t, θ_t and some coefficients, ω , to yield

Table 1. Parameterization of the model.

Preferences			Technology			
σ	β	b	α	δ	ρ	σ_ε
3	0.95	0.5	0.36	0.06	0.95	0.007

$$(c_t(\omega) - bh_t(\omega))^{-\sigma} = \beta\psi(\omega; k_t(\omega), h_t(\omega), \theta_t),$$

where $\psi(\omega; k_t(\omega), h_t(\omega), \theta_t) = \exp(P_n(\omega; \log k_t(\omega), \log h_t(\omega), \log \theta_t))$, and P_n denotes polynomial of degree n . By using the exponent of the logarithmic polynomial expansion we guarantee that the left hand side of (11) remains positive. Given $c_t(\omega)$, the next period values for the capital and habit stocks follow directly from the laws of motion (12) and (13).

3) Using the realizations of $\{\theta_t\}_{t=0}^T$, repeat the previous step in order to obtain recursively a series of the endogenous variables $\{c_t(\omega), k_{t+1}(\omega), h_{t+1}(\omega)\}_{t=0}^T$, for this particular parameterization ω .

4) Run the following non-linear regression

$$Y_t(\omega) = \exp(P_n(\xi; \log k_t(\omega), \log h_t(\omega), \log \theta_t)) + \eta_t,$$

where the role of the dependent variable $Y_t(\omega)$ is performed by the expression inside the conditional expectation in Equation (11).

5) Letting $S(\omega)$ be the result of the regression in the previous step, use an iterative procedure to find the fixed point of S , and the set of coefficients $\omega_f = S(\omega_f)$. This would provide the solution for the endogenous variables $\{c_t(\omega_f), k_{t+1}(\omega_f), h_{t+1}(\omega_f)\}_{t=0}^T$ for this particular realization of the stochastic process $\{\theta_t\}_{t=1}^T$ along with the approximated policy functions:

$$c_t(k_t, h_t, \theta_t) = bh_t + [\beta\psi(\omega_f; k_t, h_t, \theta_t)]^{-\frac{1}{\sigma}},$$

$$k_{t+1}(k_t, h_t, \theta_t) = \theta_t k_t^\alpha + (1 - \delta)k_t - bh_t - [\beta\psi(\omega_f; k_t, h_t, \theta_t)]^{-\frac{1}{\sigma}},$$

$$h_{t+1}(k_t, h_t, \theta_t) = bh_t + [\beta\psi(\omega_f; k_t, h_t, \theta_t)]^{-\frac{1}{\sigma}}.$$

Step II (Monte Carlo Integration) Our objective is to find approximations of partials at a range of points. Supposing that the point of interest is $(\bar{k}, \bar{h}, \bar{\theta})$ the algorithm proceeds as follows:

Simulate N sequences of the realizations of the stochastic process $\{\theta^n\}_{t=0}^{\bar{T}}$ of size \bar{T} with a starting value $\theta_0^n = \bar{\theta}$, for all $n = 1, \dots, N$. For a each sequence $\{\theta^n\}_{t=0}^{\bar{T}}$ simulate the series of the endogenous variables $\{k_t^n, h_t^n, c_t^n\}_{t=0}^{\bar{T}}$ using approximated policy functions, the laws of motion (12)-(13), and the corresponding initial values $k_0^n = \bar{k}, h_0^n = \bar{h}$. Using the simulated series calculate the discounted sums of the instantaneous utilities and average over N ,

$$V(\bar{k}, \bar{h}, \bar{\theta}) \simeq \frac{1}{N} \sum_{n=1}^N \sum_{t=0}^{\bar{T}} \beta^t \frac{(c_t^n - bh_t^n)^{1-\sigma}}{1-\sigma}.$$

Step III (Numerical Differentiation) To obtain $V_k(\bar{k}, \bar{h}, \bar{\theta})$ get approximations of the optimal value function at $V(k + \epsilon, h, \theta)$ and $V(k - \epsilon, h, \theta)$, where ϵ is a small positive number. Calculate the approximated value of the partial derivative using Stirling's finite difference formula:

$$\frac{\partial V(\bar{k}, \bar{h}, \bar{\theta})}{\partial k} \simeq \frac{V(\bar{k} + \epsilon, \bar{h}, \bar{\theta}) - V(\bar{k} - \epsilon, \bar{h}, \bar{\theta})}{2\epsilon}. \quad (14)$$

The partial with respect to the habit stock is obtained in a similar way. The length of the simulated series \bar{T} can be very moderate due to discounting of the future utilities. The optimal value of ϵ is both computer and problem specific.

5. Numerical Accuracy: A Comparison

This section considers the accuracy of the algorithm in the context of our example. First, we compare performance of our algorithm with the approach in [3] when such comparison is feasible. Next, we present several special cases, which isolate the contributions of different sources to the overall approximation error.

Consider the optimality conditions for the autarkic problem, written in the sequence form:

$$u_c(c_t, h_t) + \beta\lambda E_t[V_h(k_{t+1}, h_{t+1}, \theta_{t+1})] = \beta E_t[V_k(k_{t+1}, h_{t+1}, \theta_{t+1})], \quad (15)$$

$$V_k(k_t, h_t, \theta_t) = \beta E_t[V_k(k_{t+1}, h_{t+1}, \theta_{t+1})](f_k(k_t, \theta_t) + 1 - \delta), \quad (16)$$

$$V_h(k_t, h_t, \theta_t) = u_h(c_t, h_t) + \beta(1 - \lambda)E_t[V_h(k_{t+1}, h_{t+1}, \theta_{t+1})]. \quad (17)$$

Condition (17) can be used to compare our algorithm with the method in [3]. The latter requires solving the model numerically and expressing the derivatives of interest in terms of conditional expectations and functions of the equilibrium path of the model. Applying recursive substitution and the law of iterated expectations to (17) yields:

$$V_h(k_t, h_t, \theta_t) = u_h(c_t, h_t) + E_t\left[\sum_{j=1}^{\infty} \beta^j (1 - \lambda)^j u_h(c_{t+j}, h_{t+j})\right].$$

An approximation of this derivative can be obtained by parameterizing the right hand side with flexible functional forms in the state variables (k_t, h_t, θ_t) and running a non-linear regression using the simulated series from the numerical solution of the model.

Figure 1 shows the approximations of the derivative obtained using our algorithm and the approach in [3]. The approximated values of $V_h(k_t, h_t, \theta_t)$ are plotted for a range of a state variable while keeping the remaining states fixed at their deterministic steady state values. The histograms plot the sample distributions of capital and habit stock. A few observations can be made based on **Figure 1**. First, the two algorithms produce indistinguishable results when the state variables take the values which often occur in equilibrium. Second, for the points which are unlikely to occur in equilibrium, the approximations differ signifi-

cantly. To see this feature, consider the range of values of capital stock in excess of 6.5. The plots of the approximate derivatives reported in the upper panel of **Figure 1** do not coincide. Moreover, the upper tail of the histogram suggests that such values of k_t are not unlikely to happen in equilibrium. Notice, that while considering a relatively high value of k_t we kept the remaining arguments of $V_h(k_t, h_t, \theta_t)$ at their deterministic steady state values. However, the points where capital is very high while consumption (and hence the habit stock) are at the steady state level are rather unusual. This is an expected result, since Monte Carlo integration delivers good approximations only in the region of the state space which is frequently visited by the model in equilibrium.

The framework we have chosen for a worked out example embeds several well known special cases. For instance, if $\lambda = 0$, it reduces to the Brock-Mirman stochastic growth model. In this case, the analytical form of the one-period return function r , which maps the graph A of the feasibility correspondence Γ into the real numbers is known. The correspondence describing the feasibility constraints is given by

$$\Gamma(k_t, \theta_t) = [f(1-\delta)k_t, f(k_t, \theta_t) + (1-\delta)k_t],$$

and the instantaneous return function becomes

$$r : A \rightarrow \mathbb{R} \text{ given by } r(k_t, k_{t+1}, \theta_t) = u(f(k_t, \theta_t) + (1-\delta)k_t - k_{t+1}),$$

where $A = \{(k_t, k_{t+1}, \theta_t) \in \mathbb{R}_+^3 : k_{t+1} \in \Gamma(k_t, \theta_t)\}$. Hence, by virtue of the Benveniste-Scheinkman theorem the derivative of interest can be expressed as

$$V_k(k_t, \theta_t) = u'(f(k_t, \theta_t) + (1-\delta)k_t - g(k_t, \theta_t)) [f_k(k_t, \theta_t) + 1 - \delta],$$

where g is the optimal policy function for capital stock. This special case allows us to compare the simulation from our algorithm with the example where the only source of approximation errors is the approximation of the policy function g . This will allow us to isolate the contribution of the approximation errors in evaluation of the integrals and numerical differentiation to the overall approximation error of the algorithm. As shown in **Figure 1**, the approximation delivered by our algorithm is very close to the approximation which relies on the Benveniste-Scheinkman theorem. Once again, in the region of the state space which is often visited by the model in equilibrium the two approximations are virtually identical. This allows us to tentatively suggest that the main contribution to the approximation error of the algorithm comes from the approximation of the policy functions.

Our final special case compares the approximation of the derivative with its known exact solution. It is well known that for the functional forms $f(k_t, \theta_t) = \theta_t k_t^\alpha$, $u(c_t) = \log c_t$, and $\delta = 1$, the optimal policy function is defined by the simple law of motion $k_{t+1} = \alpha\beta\theta_t k_t^\alpha$. Moreover, the derivative of the value function has the following analytical solution:

$$V_k(k_t, \theta_t) = \frac{\alpha\theta_t k_t^{\alpha-1}}{(1-\alpha\beta)\theta_t k_t^\alpha} = \frac{\alpha}{1-\alpha\beta} \frac{1}{k_t}.$$

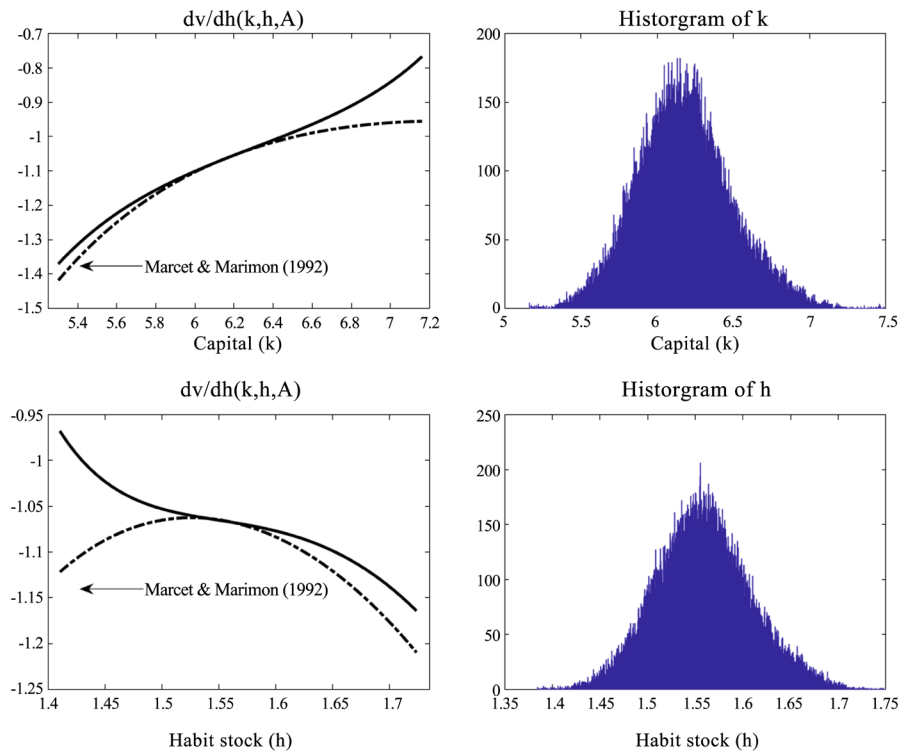


Figure 1. Comparison with the algorithm of Marcet and Marimon [3].

By replacing the approximated policy function with the known closed form solution, we can isolate the effect of the errors stemming from Monte Carlo integration and numerical differentiation on the accuracy of the approximation. **Figure 2** compares the approximated derivative obtained using the exact policy function for k_t with its analytical counterpart. The reported graphs are visually indistinguishable for the range of six standard deviations of k_t around its deterministic steady state value. The approximation errors stemming from Monte Carlo integration and numerical differentiation are of an order of 10^{-9} of the value of the derivative. This suggests that obtaining accurate approximation of the policy functions is crucial for the accuracy of the whole algorithm.

6. Concluding Remarks

This paper contributes to the literature on recursive contracts by proposing an algorithm for computing the gradient of the optimal value function using simulation-based techniques. The proposed procedure is conceptually straightforward, computationally inexpensive, and simple to implement. It allows researchers to extend existing risk-sharing models with limited commitment by including additional endogenous state variables. For example, one may extend a multi-country single-good model in [6] by including different types of capital or time-non-separable preferences. Alternatively, a multi-country multi-good model in [14] can be extended to include capital accumulation and cross-country trade in investment goods.

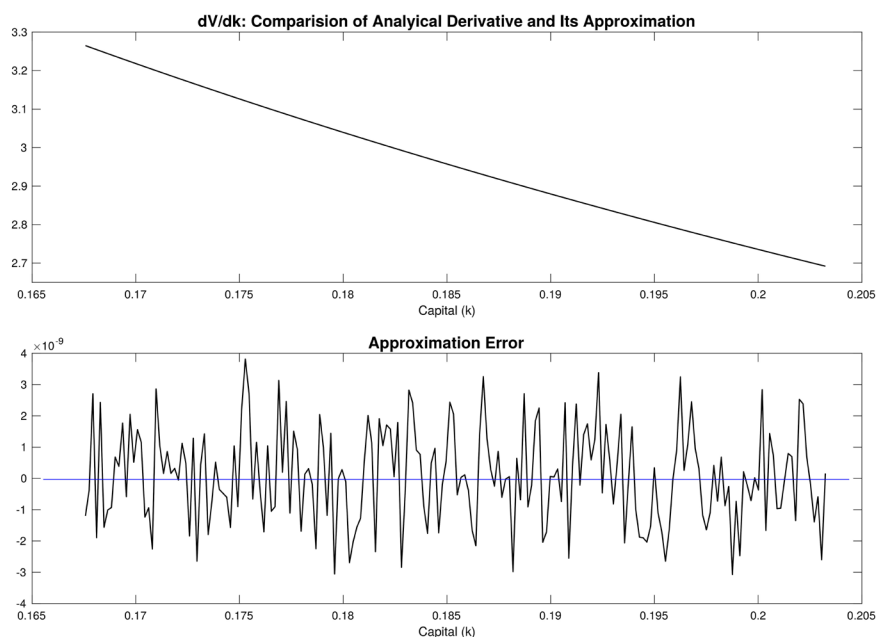


Figure 2. Comparison with a Closed Form Solution.

In terms of accuracy, the algorithm demonstrates performance comparable with Marcat and Marimon's method [3] in our benchmark example. In contrast to [3], our method is flexible enough to handle dynamic models with large numbers of state variables even when derivatives of interest cannot be expressed in terms of conditional expectations and functions of the equilibrium path of the model.

While our algorithm has wide applicability, it inherits its speed and accuracy trade-offs from the underlying numerical solution method. Our experiments suggest that obtaining accurate approximation of the policy functions is crucial for the accuracy of the whole algorithm.

An additional limitation on the algorithm's computational speed is imposed by model's simulation and Monte-Carlo integration. However, both steps can be parallelized along the lines proposed in [15] in order to reduce the computational time burden. Exploring the costs and benefits of a parallel implementation of the algorithm is left for the future research.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

- [1] Cole, H. and Kubler, F. (2012) Recursive Contracts, Lotteries and Weakly Concave Pareto Sets. *Review of Economic Dynamics*, **15**, 479-500.
- [2] Golosov, M., Tsyvinski, A. and Werquin, N. (2016) Recursive Contracts and Endogenously Incomplete Markets. In: *Handbook of Macroeconomics*, Vol. 2, Elsevier, Amsterdam, 725-841.

- <https://doi.org/10.1016/bs.hesmac.2016.03.007>
- [3] Marcet, A. and Marimon, R. (1992) Communication, Commitment, and Growth. *Journal of Economic Theory*, **58**, 219-249.
[https://doi.org/10.1016/0022-0531\(92\)90054-L](https://doi.org/10.1016/0022-0531(92)90054-L)
- [4] Benveniste, L.M. and Scheinkman, J.A. (1979) On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica*, **47**, 727-732.
<https://doi.org/10.2307/1910417>
- [5] Backus, D.K., Kehoe, P.J. and Kydland, F.E. (1992) International Real Business Cycles. *Journal of Political Economy*, **100**, 745-775.
<https://doi.org/10.1086/261838>
- [6] Kehoe, P.J. and Perri, F. (2002) International Business Cycles with Endogenous Incomplete Markets. *Econometrica*, **70**, 907-928.
<https://doi.org/10.1111/1468-0262.00314>
- [7] Dmitriev, A. and Roberts, I. (2012) International Business Cycles with Complete Markets. *Journal of Economic Dynamics and Control*, **36**, 862-875.
<https://doi.org/10.1016/j.jedc.2011.12.006>
- [8] Dmitriev, A. and Krznar, I. (2012) Habit Persistence and International Comovements. *Macroeconomic Dynamics*, **16**, 312-330.
<https://doi.org/10.1017/S1365100510000957>
- [9] Chen, X. and Ludvigson, S.C. (2009) Land of Addicts? An Empirical Investigation of Habit-Based Asset Pricing Models. *Journal of Applied Econometrics*, **24**, 1057-1093.
- [10] Judd, K.L. (1998) Numerical Methods in Economics. The MIT Press, Cambridge, MA.
- [11] Mathews, J.H. and Fink, K.K. (2004) Numerical Methods Using Matlab. 4th Edition, Prentice-Hall, Upper Saddle River, NJ.
- [12] Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P. (1992) Numerical Recipes in C: The Art of Scientific Computing. 2nd Edition, Cambridge University Press, Cambridge.
- [13] den Haan, W. and Marcet, A. (1990) Solving the Stochastic Growth Model by Parameterizing Expectations. *Journal of Business & Economic Statistics*, **8**, 31-34.
- [14] Bodenstein, M. (2008) International Asset Markets and Real Exchange Rate Volatility. *Review of Economic Dynamics*, **11**, 688-705.
<https://doi.org/10.1016/j.red.2007.12.003>
- [15] Creel, M. (2008) Using Parallelization to Solve a Macroeconomic Model: A Parallel Parameterized Expectations Algorithm. *Computational Economics*, **32**, 343-352.
<https://doi.org/10.1007/s10614-008-9142-6>