

A Spatial Interpolation Framework for Efficient Valuation of Large Portfolios of Variable Annuities

Seyed Amir Hejazi^a, Kenneth R. Jackson^a, Guojun Gan^b

^a*Department of Computer Science, University of Toronto, Toronto, ON, M5S 3G4, Canada*

^b*Department of Mathematics, University of Connecticut, Storrs, Connecticut, 06269-3009, USA*

Abstract

Variable Annuity (VA) products expose insurance companies to considerable risk because of the guarantees they provide to buyers of these products. Managing and hedging these risks requires insurers to find the value of key risk metrics for a large portfolio of VA products. In practice, many companies rely on nested Monte Carlo (MC) simulations to find key risk metrics. MC simulations are computationally demanding, forcing insurance companies to invest hundreds of thousands of dollars in computational infrastructure per year. Moreover, existing academic methodologies are focused on fair valuation of a single VA contract, exploiting ideas in option theory and regression. In most cases, the computational complexity of these methods surpasses the computational requirements of MC simulations. Therefore, academic methodologies cannot scale well to large portfolios of VA contracts. In this paper, we present a framework for valuing such portfolios based on spatial interpolation. We provide a comprehensive study of this framework with existing interpolation schemes. Our numerical results show superior performance, in terms of both computational efficiency and accuracy, for these methods compared to nested MC simulations. We also present insights into the challenge of finding an effective interpolation scheme in this framework, and suggest guidelines that help us build a fully automated scheme that is efficient and accurate.

Keywords: Variable annuity, Spatial interpolation, Kriging, Inverse

Email addresses: amir@cs.toronto.edu (Seyed Amir Hejazi), krj@cs.toronto.edu (Kenneth R. Jackson), Guojun.Gan@uconn.edu (Guojun Gan)

1. Introduction

Variable annuities are unit-linked products that are wrapped with a life insurance contract. These products allow a policyholder to invest into pre-defined sub-accounts set up by the insurance company. Sub-account funds are invested in bonds, the money market, stocks and other financial products. An insurer offers different types of sub-accounts that are tailored to the appetite of policyholders with different level of tolerance for risk. The investment can be made via a lump-sum payment or a series of investment purchases. In return, the insurance company offers tax sheltered growth and guarantees that protect the policyholder in a bear market (Chi and Lin, 2012).

Upon entering into a contract, the policyholder is given two accounts: 1) keeps track of the performance of investments in the sub-accounts 2) keeps track of the amount of guarantee provided by the insurance company. The value of former account is called the account value and the value of latter account is called the benefit base. During a period called the accumulation phase, the policyholder accumulates assets on his investments in sub-accounts and the value of his benefit base appreciates by contractually agreed roll ups, ratchets and resets without taxation. At the end of the accumulation phase, the benefit base is locked in and the insurer guarantees to return at least the benefit base as a lump sum payment or as a stream of payments during a period called the withdrawal phase.

The most prevalent of the guarantees are the Guaranteed Minimum Death Benefit (GMDB), the Guaranteed Minimum Withdrawal Benefit (GMWB), the Guaranteed Minimum Income Benefit (GMIB), and the Guaranteed Minimum Accumulation Benefit (GMAB). The GMDB guarantees a specified lump sum payment on death regardless of the performance of the underlying account. The most basic guarantee offered now is the return of benefit base adjusted for any partial withdrawals. The GMWB guarantees the ability to partially withdraw up to a pre-determined percentage (called the withdrawal rate) of the benefit base for a specified number of years. The decision to withdraw is made annually and the maximum amount of withdrawal is a function of the age of the policyholder. The GMIB guarantees a stream of income for life contingent on the survival of the policyholder, and the GMAB guarantees a lump sum payment on maturity of the contract regardless of the performance of the underlying funds. For further details, see (TGA, 2013).

Embedded guarantees are the key selling feature of VA products. These guarantees have allowed insurance companies to sell trillions of dollars worth of these products worldwide, in 2010 alone (IRI, 2011). As a result, major insurance companies are now managing large portfolios of VA contracts, each with hundreds of thousands of contracts.

Although the embedded guarantees are attractive features to the buyer of VA products, they expose the insurers to substantial risk (e.g., market and behavioral risk). Because of that, major insurance companies have started risk management and hedging (Boyle and Hardy, 1997; Hardy, 2003) programs, especially after the market crash of 2008, to reduce their exposures. An integral part of a risk management program is finding the value of key statistical risk indicators, e.g., the Greeks (Hull, 2006), and the Solvency Capital Requirement (SCR) (Bauer et al., 2012), on daily, monthly and quarterly bases.

Most of the academic research to date has focused on fair valuation of individual VA contracts (Coleman and Patron, 2006; Milevsky and Salisbury, 2006; Boyle and Tian, 2008; Lin et al., 2008; Belanger et al., 2009; Gerber et al., 2012; Chen et al., 2008; Dai et al., 2008; d’Halluin et al., 2005; Azimzadeh and Forsyth, 2013; Chen and Forsyth, 2008; Huang and Forsyth, 2011; Moenig and Baue, 2011; Ulm, 2006). Most of the methodologies developed in these research papers are based on ideas from option pricing theory, and are tailored to the type of studied VA. In addition, almost all of the proposed schemes are computationally expensive and the results they provide for a VA contract cannot be re-used for another VA contract, even of similar type. Each VA contract is unique in terms of its key attributes, i.e., age, gender, account value, guaranteed value, maturity of contract, fund type, etc. Hence, VA portfolios are non-homogeneous pools of VA contracts, and academic methodologies cannot scale well to be used to calculate key risk statistics of large VA portfolios.

Although the nature of the guarantees in the VA products makes them path dependent, in practice, insurance companies relax the assumptions on guarantees and rely heavily on stochastic simulations to value these products and manage their risks. In particular, nested MC simulations are the industry standard methodology in determining key risk metrics (Bauer et al., 2012; Reynolds and Man, 2008). Nested simulations consist of outer loops that span the space of key market variables (risk factors), and inner loops that project the liability of each VA contract along many simulated risk-neutral paths (Fox, 2013). As explained in Section 4, MC simulations are computa-

tionally demanding, forcing insurance companies to look for ways to reduce the computational load of MC simulations.

In this paper, we focus on the issue of efficient approximation of the Greeks for large portfolios of VA products. In particular, we provide an extensive study of a framework based on metamodeling that can approximate the Greeks for large portfolios of VAs in a fast and accurate way. The rest of this paper is organized as follows. Section 2 reviews existing efforts to reduce the computational requirements of MC simulations. Section 3 introduces the spatial interpolation framework that offers fast and accurate valuation of large VA portfolios. We discuss various ways that this method can be employed with existing interpolation schemes to approximate the Greeks of VA portfolios. Section 4 provides insights into the performance of the proposed framework via extensive simulations. Finally, Section 5 concludes the paper.

2. Review of Existing Methods

VAs can be considered to be a type of exotic market instrument (Hull, 2006). Hence, in this section, we provide a summary of the main existing approaches to the more general problem of valuing large portfolios of exotic market instruments. As we discussed earlier, MC simulations are the industry standard approach to value path dependent products. The problem with MC simulations is that they are computationally demanding and do not scale well to large portfolios of exotic instruments. The existing methods try to alleviate the computational burden of MC simulations by approximating the surface that defines a mapping between the key risk metrics of interest and the key economic factors. Most of these approaches are based on ideas in regression theory and require a time consuming calibration to find the regression coefficients.

A well-studied approach in the literature is the method of replicating portfolios. In a replicating portfolio, the main idea is to regress the cash flow of the input portfolio, at each time step, against the cash flow of a portfolio of well-formulated financial instruments, in particular, vanilla derivatives. The financial instruments that are used in regression often have closed form formulas for liability values. Hence, simplifying the calculation of cash flows and subsequently reducing the cost of finding regression coefficients. The problem of determining a replicating portfolio is often formulated as a minimization problem with respect to a norm. Depending on the choice of norm,

quadratic programming (Daul and Vidal, 2009; Oechslin et al., 2007) and linear programming (Dembo and Rosen, 1999) approaches have been proposed to solve the optimization problem. However, constructing the replicating portfolio via either method is time consuming because it requires projecting the cash flow at each time step. Moreover, depending on the desired level of accuracy, the number of regression variables can grow polynomially with the number of attributes of the instruments in the input portfolio. Another major issue with replicating portfolios is the difficulty in incorporating the actuarial risk factors. The usual practice in the literature is to assume that these values follow their expected value.

Another important regression approach is the method of Least Squares Monte Carlo (LSMC). In LSMC, the liability value is regressed on key economic factors (Cathcart and Morrison, 2009). The common practice is to choose powers of key economic factors as basis functions and approximate the liability as a linear combination of the basis functions (Carriere, 1996; Longstaff and Schwartz, 2001). Hence, the Greeks of the input portfolio can be calculated as the sum of the estimated Greeks for each VA in the input portfolio. In order to find regression coefficients, LSMC uses poor MC estimates of liability at each point in the space of regression variables. Cathcart and Morrison (Cathcart and Morrison, 2009) provide examples in which the number of MC projection paths to estimate liability at each point in the space is reduced to one and yet the regression function provides fairly accurate estimates. However, to achieve a reasonable accuracy, LSMC usually requires many sample points for numeric explanatory variables that can assume values in a large interval. Each VA product has several numeric attributes, some of which are unique to the type of VA contract. Therefore, the LSMC method incurs significant computational costs to accurately regress the liability of contracts in a portfolio consisting of different types of VA products.

Recently, a new spatial functional data analysis approach (Gan and Lin, 2015; Gan, 2013a,b) has been proposed that addresses the computational complexity of nested simulations by reducing the number of VA contracts included in the MC simulations. The proposed methods first select a small set of representative VA policies, using various data clustering (Gan et al., 2007) or sampling methods, and price only these representative policies via MC simulations. The representative contracts and their Greeks are then fed as training samples to a machine learning algorithm (Bishop, 2006) called Kriging (Cressie, 1993), which then estimates the Greeks of the whole portfolio. In the rest of the paper, we provide a study of the more general framework of

spatial interpolation, including Kriging methods, and provide more insights into why spatial interpolation can be much more efficient and accurate than other approaches in the literature. In this paper, we use the term interpolation in the general context of estimating the values at unknown locations using the known data at a finite number of points. In this context, an interpolation method that predicts the values at points for which the data values are known to be the known values is called an *exact* interpolator. Interpolation methods that do not satisfy this constraint are called *inexact* interpolation methods (Burrough et al., 1998).

3. Spatial Interpolation Framework

The proposed methods in (Gan and Lin, 2015; Gan, 2013a,b) can be categorized under the general framework of spatial interpolation. Spatial interpolation is the procedure of estimating the value of data at unknown locations in space given the observations at sampled locations (Burrough et al., 1998). As the definition suggests, spatial interpolation requires a sampling method to collection information about the surface of interest and an interpolation method that uses the collected information to estimate the value of the surface at unknown locations. As discussed in (Gan and Lin, 2015; Gan, 2013a,b), the choice of sampling method and interpolation method can noticeably impact the quality of the interpolation. In this paper, we choose to focus on the latter, and leave for another paper a discussion of the choice of an appropriate sampling method.

In the functional data analysis literature, there exist two main classes of interpolation methods (Burrough et al., 1998):

- **Deterministic Interpolation:** Creates surfaces from measured points on the basis of either similarity or degree of smoothness.
- **Stochastic Interpolation:** Utilizes statistical properties of measured points, such as auto-correlation amongst measured points, to create the surface.

In what follows, we study three (one stochastic, and two deterministic) of the most prominent of these interpolation techniques— Kriging, Inverse Distance Weighting (IDW) and Radial Basis Function (RBF)— in the context of our problem of interest. In particular, we study these interpolation techniques to estimate the delta value for a large portfolio of VA products.

Although our study focuses on estimation of the delta value, the framework is general and can be applied to estimate other Greeks as well. We compare the performance of these methods in terms of computational complexity and accuracy at the micro (contract) level and at the macro (portfolio) level.

Although (Gan and Lin, 2015; Gan, 2013a) provide some insights into the performance of the Kriging interpolation methods, we provide further insights into the efficiency and accuracy of Kriging based methods in comparison to other interpolation techniques. Moreover, we shed some light on how Kriging achieves its documented performance and discuss some issues regarding the choice of variogram model and distance function.

3.1. Sampling Method

In this paper, we focus on studying synthetic portfolios that are generated uniformly at random in the space of selected variable annuities. In (Gan, 2013b), the Latin Hypercube Sampling (LHS) method (McKay et al., 1979) is used to select representative contracts. LHS provides a uniform coverage of the space including the boundary VA contracts. The results of (Gan, 2013b) indicate that LHS increases the accuracy of the estimation compared to other sampling methods. In order to preserve the properties of LHS, we select our representative contracts by dividing the range of each numeric attribute of a VA contract into almost equal length subintervals, selecting the end points of resulting subintervals and producing synthetic contracts from all combinations of these end points and choices of categorical attributes.

3.2. Kriging

Kriging is a stochastic interpolator that gives the best linear unbiased estimation of interpolated values assuming a Gaussian process model with prior covariance (Matheron, 1963; Krige, 1951). Various Kriging methods (i.e., Simple Kriging, Ordinary Kriging, Universal Kriging, etc.) have been developed based on assumptions about the model. In our experiments, we didn't find any significant advantages in choosing a particular Kriging method. Therefore, for the sake of simplicity of analysis, and based on the results of (Gan and Lin, 2015), we choose to study ordinary Kriging in this paper.

Assume $Z(x)$ represents the delta value of a VA contract represented in space by the point x . Let $Z(x_1), Z(x_2), \dots, Z(x_n)$ represent the observed delta values at locations x_1, x_2, \dots, x_n . Ordinary Kriging tries to find the best, in the Mean Squared Error (MSE) sense, unbiased linear estimator

$\hat{Z}(x) = \sum_{i=1}^n \omega_i Z(x_i)$ of $Z(x)$ by solving the following system of linear equations to find the w_i s.

$$\begin{bmatrix} \gamma(D(x_1, x_1)) & \gamma(D(x_1, x_2)) & \dots & \gamma(D(x_1, x_n)) & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \gamma(D(x_n, x_1)) & \gamma(D(x_n, x_2)) & \dots & \gamma(D(x_n, x_n)) & 1 \\ 1 & 1 & \dots & 1 & 0 \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_n \\ \lambda \end{bmatrix} = \begin{bmatrix} \gamma(D(x_1, x)) \\ \dots \\ \gamma(D(x_n, x)) \\ 1 \end{bmatrix} \quad (1)$$

where λ is the Lagrange multiplier (Boyd and Vandenberghe, 2004), $\gamma(\cdot)$ is the semi-variogram function, to be discussed shortly, and $D(\cdot, \cdot)$ is a distance function that measures the distance between two points in the space of VA contracts. The last row enforces the following constraint to allow an unbiased estimation of $Z(x)$.

$$\sum_{i=1}^n w_i = 1 \quad (2)$$

In this formulation of the Kriging problem, the system of linear equations (1) should be solved once for each VA policy (point in space). Solving a system of linear equations, with standard methods, takes $\Theta(n^3)^1$ time. Hence, estimating the delta value for a portfolio of N VA contracts by summing the estimated delta value of each contract requires $\Theta(N \times n^3)$ time which is computationally expensive. Because of this, Kriging methods are inefficient in finding a granular view of the portfolio. However, if we are only interested in the Greeks of the portfolio, and not the Greeks of each individual policy, we can follow the approach of (Gan and Lin, 2015; Gan, 2013a,b) and use the linearity of the systems of linear equations to sum them across the portfolio in $\Theta(n \times N)$ and to solve only the resulting system of linear equations in time proportional to n^3 . Hence estimating the delta of a portfolio requires $\Theta(n^3 + n \times N)$ time. To sum the systems of linear equations, we sum the corresponding weights and Lagrange multipliers on the left hand side of the equations and sum the corresponding terms, i.e., $\gamma(D(x_i, x))$, $i = 1, 2, \dots, n$, and constants, on the right hand side of the equations.

¹ $f(x) = \Theta(g(x))$ means that there exists positive numbers c_1 , c_2 , and M such that $\forall x > M : c_1 g(x) \leq f(x) \leq c_2 g(x)$.

3.2.1. Variogram

Kriging assumes the Gaussian process $Z(x)$ is second order stationary, i.e., the covariance of the Gaussian process in two locations is a function of distance between the two locations. Assuming a zero mean, the Gaussian process covariance function can be defined in terms of a variogram function $2\gamma(h)$

$$\begin{aligned} Cov_Z(x+h, x) &= E[Z(x+h)Z(x)] \\ &= \frac{1}{2}E\left[Z^2(x+h) + Z^2(x) - \left(Z(x+h) - Z(x)\right)^2\right] \\ &= Var(Z) - \frac{1}{2}(2\gamma(h)) \end{aligned} \quad (3)$$

In practice, for a set of sample points $x_i, 1 \leq i \leq n$, the variogram can be estimated as

$$2\hat{\gamma}(h) = \frac{1}{N(h)} \sum_{i=1}^{N(h)} (Z(x_i+h) - Z(x_i))^2 \quad (4)$$

where $N(h)$ is the number of pairs in the sample separated by a distance h from each other. The function $2\hat{\gamma}(h)$ is often called the *empirical variogram*.

Because of the noise in measurements, the estimated empirical variogram may not represent a valid variogram function. Since methods like Kriging require valid variograms at every distance h , empirical variograms are often approximated by model functions ensuring the validity of the variogram (Chiles and Delfiner, 1999). Variogram models are usually described in terms of three important variables:

- **Nugget (n):** The height of the discontinuity jump at the origin.
- **Sill (s):** The Limit of the variogram as the lag distance h approaches infinity.
- **Range (r):** The distance at which the difference of the variogram from the sill becomes negligible.

Figure 1 shows an example of an empirical variogram and the model variogram. In our study, we choose to focus on the following three prominent variogram models (Chiles and Delfiner, 1999; Cressie, 1993):

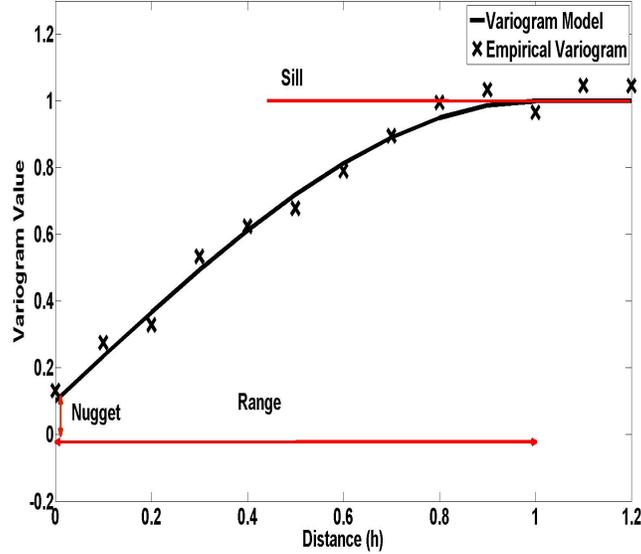


Figure 1: Example of a variogram.

- Exponential Variogram

$$\gamma(h) = (s - n) \left(1 - \exp \left(- \frac{h}{(ra)} \right) \right) + n \mathbf{1}_{(0, \infty)}(h)$$

- Spherical Variogram

$$\gamma(h) = (s - n) \left(\left(\frac{3h}{2r} - \frac{h^3}{2r^3} \right) \mathbf{1}_{(0, r)}(h) + \mathbf{1}_{[r, \infty)}(h) \right) + n \mathbf{1}_{(0, \infty)}(h)$$

- Gaussian Variogram

$$\gamma(h) = (s - n) \left(1 - \exp \left(- \frac{h^2}{r^2 a} \right) \right) + n \mathbf{1}_{(0, \infty)}(h)$$

In Exponential and Gaussian variogram models, a is a free parameter that is chosen so that the variogram better fits the data.

3.3. Inverse Distance Weighting

Inverse Distance Weighting (IDW) is a deterministic method that estimates the value at an unknown position x as a weighted average of values at

known positions x_1, \dots, x_n . Assuming the delta values $Z(x_1), Z(x_2), \dots, Z(x_n)$ of representative VAs x_1, x_2, \dots, x_n , we can estimate the delta value $Z(x)$ of a VA at a point x as

$$\hat{Z}(x) = \begin{cases} \frac{\sum_{i=1}^n w_i(x)Z(x_i)}{\sum_{i=1}^n w_i(x)} & \forall i : D(x, x_i) \neq 0 \\ Z(x_i) & \exists i : D(x, x_i) = 0 \end{cases} \quad (5)$$

where $w_i(x) = D(x, x_i)^{-p}$, and $D(\cdot, \cdot)$ is a distance function (Shepard, 1968). The parameter p in $w_i(x)$ is a positive real number called the ‘‘power parameter’’. The choice of power parameter depends on the distribution of samples and the maximum distance over which an individual sample is allowed to influence the surrounding points. Greater values of p assign greater influence to values closest to the interpolating point. The choice of the power parameter also influences the smoothness of the interpolator by changing the influence radius of sample points.

In comparison to Kriging, IDW requires only $\Theta(n)$ operations to estimate the delta value of each new VA contract using the delta values of n representative contracts. Assuming a portfolio of N VA contracts, we can estimate the delta value of the portfolio by summing the estimated delta value of contracts in time proportional to $n \times N$. Hence, we expect IDW to be faster than Kriging to estimate the delta value of the portfolio. The difference in speed is more apparent if we want a more granular view of the portfolio. In other words, if we are interested in the estimated delta value of each VA contract in the portfolio, Kriging is much slower than IDW. We provide further insights into this matter in Section 4.

3.4. Radial Basis Functions

In the RBF method, we approximate the delta value of a VA contract x as a weighted linear combination of radial functions centered at representative contracts x_1, x_2, \dots, x_n :

$$\hat{Z}(x) = \sum_{i=1}^n w_i \Phi(\|x - x_i\|) \quad (6)$$

where $\|\cdot\|$ is a norm, usually chosen to be Euclidean distance.

In RBF interpolation, the weights are chosen so that RBF is exact at the $x_i, 1 \leq i \leq n$, points. In other words, given the values $Z(x_1), \dots, Z(x_n)$ at points x_1, \dots, x_n , the following linear set of equations is solved for w_i :

$$\begin{bmatrix} \Phi(\|x_1 - x_1\|) & \dots & \Phi(\|x_1 - x_n\|) \\ \vdots & \ddots & \vdots \\ \Phi(\|x_n - x_1\|) & \dots & \Phi(\|x_n - x_n\|) \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} = \begin{bmatrix} Z(x_1) \\ \vdots \\ Z(x_n) \end{bmatrix} \quad (7)$$

In our research, we chose the following commonly used radial basis functions

- Gaussian

$$\Phi(x) = \exp(-\epsilon x^2) \quad (8)$$

- Multi-Quadratic

$$\Phi(x) = \sqrt{1 + (\epsilon x)^2} \quad (9)$$

These two functions represent two classes of radial basis functions: 1) the class in which the value of the radial function increases with the distance from its center, 2) the class in which the value of radial function decreases with the distance from its center. Although the latter class of RBF functions, which is represented by Gaussian radial function in our study, seems more suitable for our application of interest, for the sake of completeness, we chose to experiment with the former class as well in our study. In both of the above-mentioned functions, ϵ is a free parameter that defines the significance of known points on the value of their neighbor points.

Similar to IDW, RBF interpolation has a running time that is proportional to n for the delta value estimation of each VA contract, and a running time of $\Theta(n \times N)$ to estimate the delta value of a portfolio of N VA contracts. But in addition we need extra $\Theta(n^3)$ time to solve (7). Hence, in total, the computational complexity of RBF interpolation to estimate the delta value of a portfolio is $\Theta(n \times N + n^3)$. Similar to IDW, the RBF interpolation can provide us more granularity in a faster time than the Kriging method.

4. Numerical Experiments

In this section, we present numerical results on the performance of each interpolation method in the context of the proposed framework. In all of our experiments, our goal is to estimate the delta value of a synthetic portfolio of 100,000 VA contracts which are chosen uniformly at random from the space defined by attributes listed in Table 1. The range of attributes are similar to the ones reported in (Gan and Lin, 2015; Gan, 2013a) which allows us to

Attribute	Value
Guarantee Type	{GMDB, GMDB + GMWB}
Gender	{Male, Female}
Age	{20, 21, ..., 60}
Account Value	[1e4, 5e5]
Guarantee Value	[0.5e4, 6e5]
Withdrawal Rate	{0.04, 0.05, 0.06, 0.07, 0.08}
Maturity	{10, 11, ..., 25}

Table 1: GMDB and GMWB attributes and their respective ranges of values.

fairly compare our results with the reported findings in (Gan and Lin, 2015; Gan, 2013a). However, for the sake of generality, we allow VA contracts, with guarantee values that are not equal to the account value. Moreover, for VA contracts with a GMWB rider, we set the guaranteed death benefit value to be equal to the guaranteed withdrawal benefit.

In our experiments, we use the framework of (Bauer et al., 2008) to value each VA contract, and assume the output of a MC simulation with 10,000 inner loop scenarios as the actual value of the contract. Inner loop scenarios are generated assuming a simple log-normal distribution model (Hull, 2006) with a risk free rate of return of $\mu = 3\%$, and volatility of $\sigma = 20\%$. Our mortality rates follow the 1996 IAM mortality tables provided by the Society of Actuaries.

In the training (calibration) stage of our proposed framework, we use MC simulations with 10,000 inner loop scenarios to find the delta value of our representative contracts. The reason behind our choice is that, when fewer inner loop scenarios are used, e.g. 1000 as used in (Gan, 2013a), we observed a noticeable difference between the computed delta value from successive runs. The observed difference when 1000 is used can be as big as 5%.

4.1. Performance

In this set of experiments, our objective is to provide a fair comparison of accuracy and computational efficiency of each proposed estimation method when the k-prototype distance function of (Gan, 2013a) is used. Since we allow the guaranteed value of VAs in the synthetic portfolio to be different than their account value, we have modified the distance function to the following.

	Experiment 1
Guarantee Type	{GMDB, GMDB + GMWB}
Gender	{Male, Female}
Age	{20, 30, 40, 50, 60}
Account Value	{1e4, 1.25e5, 2.5e5, 3.75e5, 5e5}
Guarantee Value	{0.5e4, 3e5, 6e5}
Withdrawal Rate	{0.04, 0.08}
Maturity	{10, 15, 20, 25}

Table 2: Attribute values from which representative contracts are generated for experiments.

$$D(\mathbf{x}, \mathbf{y}, \gamma) = \sqrt{\sum_{h \in N} \left(\frac{x_h - y_h}{\max_h - \min_h} \right)^2 + \gamma \sum_{h \in C} \delta(x_h, y_h)} \quad (10)$$

where $N = \{\text{AV, GD, GW, maturity, age, withdrawal rate}\}$ is the set of numerical values and $C = \{\text{gender, rider}\}$ is the set of categorical values.

Similar to (Gan and Lin, 2015; Gan, 2013a), we choose $\gamma = 1$. Moreover, we form the set of representative contracts, via the sampling method of Section 3.1, from all combinations of end points presented in Table 2. Because of the constraints on the guaranteed values, some of the entries are duplicate, which we remove to obtain a sample of size 1800.

In order to be thorough in our experiments and comprehensive in our analysis, we present the results for all variants of Kriging, IDW, and RBF methods. For Kriging, we choose to experiment with all three major variogram models, i.e., spherical, exponential, and gaussian. For IDW, we choose to experiment with different choices of the power parameter to see the effect of this free parameter on the accuracy of results. For RBF, we study two of the most popular radial functions, Gaussian and multi-quadratic, and for each type of radial function, we experimented with two choices for the free parameter ϵ . In Table 3, the relative error in estimation of the delta value of the portfolio is presented. The relative error for method m is calculated as follows.

$$\text{Err}_m = \frac{\Delta_m - \Delta_{MC}}{|\Delta_{MC}|} \quad (11)$$

where Δ_{MC} is the estimated delta value of the portfolio computed by MC

Method	Relative Error (%)
Kriging (Spherical)	-0.03
Kriging (Exponential)	-1.61
Kriging (Gaussian)	< -500
IDW ($p = 1$)	9.11
IDW ($p = 10$)	13.12
IDW ($p = 100$)	11.99
RBF (Gaussian, $\epsilon = 1$)	-1.79
RBF (Gaussian, $\epsilon = 10$)	37.89
RBF(Multi-Quad, $\epsilon = 1$)	-71.62
RBF(Multi-Quad, $\epsilon = 10$)	-10.86

Table 3: Relative error in estimation of delta value via each method.

simulations and Δ_m is the estimate delta value of the portfolio computed by method m . While two of the Kriging methods provide accurate estimates, the accuracy of IDW, and multi-quadratic RBF methods is moderate. One interesting observation is that the choice of variogram model has substantial impact on the accuracy of the method and it confirms the result of (Gan, 2013a) that the spherical method provides the best accuracy. Another interesting observation is the effect of the free parameters p and ϵ on the accuracy of the IDW and RBF methods. The results suggest that the effective use of either method requires a careful tuning of these free parameters.

Table 4 presents the running time of each algorithm in two scenarios: 1) estimating the delta value of the portfolio only 2) estimating the delta value of each VA policy in the portfolio and summing them to get the delta value of the portfolio. While the former does not provide a granular view of the portfolio, the latter gives a more refined estimation process and allows for deeper analysis and insights. Note that the times in Table 4 represent only the time that it took to estimate the values once we knew the delta values of the representative contracts. To get the total simulation time, add 187 seconds to these times, which is the time that it took to estimate the delta value of representative contracts via MC simulations. The results show the superiority of the proposed framework over MC simulation (speed up $> 15\times$) except when Kriging is used for per policy estimation of delta. Because the IDW and RBF methods by definition require the estimation of the delta of each policy and sum the estimations to get the delta value of the portfolio,

Method	Portfolio	Per Policy
MC	10617	10617
Kriging (Spherical)	312	> 320000
Kriging (Exponential)	333	> 320000
Kriging (Guassian)	383	> 320000
IDW (P = 1)	285	286
IDW (P = 10)	288	287
IDW (P = 100)	287	301
RBF (Gaussian, $\epsilon = 1$)	295	306
RBF (Gaussian, $\epsilon = 10$)	294	315
RBF(Multi-Quad, $\epsilon = 1$)	289	289
RBF(Multi-Quad, $\epsilon = 10$)	297	292

Table 4: simulation time of each method to estimate the delta value. All times are in seconds.

we can see that simulation times for these methods are approximately equal in the two presented scenarios. Moreover, these methods are more efficient than the Kriging method, which confirms our analysis in Section 3.

4.2. Accuracy

The accuracy results of Table 3 may misleadingly suggest that the Kriging method with the Spherical variogram model is always capable of providing very accurate interpolations. In the experiments of this section, we provide results on the accuracy of different methods that contradicts this hypothesis.

For our experiments in this section, we replicated the experiments of Section 4.1 with sets of representative contracts that are produced from the set of representative contracts in Section 4.1 by removing 100, 200, 400, 600 and 800 VA contracts at random. Table 5 presents the mean and the standard deviation of the relative error, in estimation of the delta value of the VA portfolio, for each method in these experiments. The results of Table 5 show high variance values for the accuracy of the Kriging methods, which contradicts our hypothesis. Another interesting observation is that the IDW methods and the RBF methods with a Gaussian kernel, in comparison to the Kriging methods, have a lower variance value for the relative error.

Method	Mean (%)	STD (%)
Kriging (Spherical)	0.47	1.76
Kriging (Exponential)	-0.58	2.19
Kriging (Gaussian)	1109.02	3289.81
IDW ($p = 1$)	9.14	1.75
IDW ($p = 10$)	13.14	0.42
IDW ($p = 100$)	12.06	0.23
RBF (Gaussian, $\epsilon = 1$)	-1.78	0.48
RBF (Gaussian, $\epsilon = 10$)	38.87	1.42
RBF(Multi-Quad, $\epsilon = 1$)	-58.65	16.84
RBF(Multi-Quad, $\epsilon = 10$)	-9.15	3.56

Table 5: Mean and standard deviation of relative error in estimation of delta value via each method.

4.3. Distance Function

A key element in the definition of each estimation method is the choice of a distance function. While the RBF method requires the choice of a proper distance function, Kriging and IDW can work with any choice of distance function. A proper distance function satisfies non-negativity, identity of indiscernibles, symmetry and the triangle inequality (Stein and Shakarchi, 2009). We call any function that has the non-negativity and a subset of other aforementioned properties a distance function. In this set of experiments, we investigate the importance of the choice of distance function on the accuracy of estimation for each interpolation method.

To achieve this goal, we conduct two sets of experiments. In the first set of experiments, we study the effect of the γ variable in (10) by reducing the value of γ from 1 to 0.05. γ determines the relative importance of the categorical attributes compared to the numerical attributes, which has not been studied previously. In the second set of experiments, we use the following distance function in our method with $\gamma = 1$.

$$\begin{aligned}
D(\mathbf{x}, \mathbf{y}, \gamma) &= \sqrt{f(x_{\text{age}}, y_{\text{age}})g_{\text{age}}(\mathbf{x}, \mathbf{y}) + \sum_{h \in N} g_h(\mathbf{x}, \mathbf{y}) + \gamma \sum_{h \in C} \delta(x_h, y_h)} \\
f(x_{\text{age}}, y_{\text{age}}) &= \exp\left(\frac{x_{\text{age}} + y_{\text{age}}}{2} - M\right) \\
g_h(\mathbf{x}, \mathbf{y}) &= (\exp(-r_x)x_h - \exp(-r_y)y_h)^2
\end{aligned} \tag{12}$$

where $C = \{\text{gender, rider}\}$, $N = \{\text{maturity, withdrawal rate}\}$, $r = \frac{AV}{GD}$ and M is the maximum age in the portfolio.

If we view the embedded guarantees in the VA contracts as options that a policyholder can choose to exercise, the ratio r represents the moneyness of that option. If $r \gg 1$, then the account value is enough to cover the amount of guaranteed value. However, if $r \ll 1$, the account value is insufficient to cover the guaranteed value and the insurer has a potential liability. Hence, in estimating the delta value for a VA contract with $r \gg 1$, the delta value is close to zero and the choice of representative contract(s) should not affect the outcome of the estimation as long as the selected representative contract(s) have $r \gg 1$. The choice of function $g(\cdot, \cdot)$ in (12) captures the aforementioned idea. In addition, as the age of the policyholder increases their mortality rate also increases (consult the data of 1996 I AM mortality table). Hence, the liability and delta value of similar contracts which differ only in the age of the policyholder increases with age. Because of this, more emphasis should be placed on estimating the delta value for senior policyholders, which is the motivation behind the introduction of the function $f(\cdot, \cdot)$ in (12).

Table 6 presents the accuracy of our estimation by each method in both experiments. In experiment one, the choice of $\gamma = 0.05$, in general, has improved the accuracy of most interpolation schemes. Kriging interpolation with a Spherical variogram, the IDW method with $P = 10$, and the RBF method with Gaussian kernel and $\epsilon = 1$ are the only schemes for which the accuracy deteriorated. In experiment two, the Kriging and RBF methods encounter singularities with (12); however, the choice of (12) has improved the accuracy of the IDW methods. In general, it seems that the choice of distance function and free parameters plays a key role in the accuracy of the interpolation schemes.

4.4. Variogram

As mentioned in Section 3.2, Kriging methods work with variogram models. The choice of variogram model is dictated by its closeness to the empirical

Method	Relative Error (%)	
	Experiment 1	Experiment2
Kriging (Spherical)	1.94	*
Kriging (Exponential)	-0.37	*
Kriging (Guassian)	< -500	*
IDW (p = 1)	8.97	-4.87
IDW (p = 10)	13.21	3.90
IDW (p = 100)	11.99	2.32
RBF (Gaussian, $\epsilon = 1$)	-2.56	*
RBF (Gaussian, $\epsilon = 10$)	37.89	*
RBF(Multi-Quad, $\epsilon = 1$)	-35.74	*
RBF(Multi-Quad, $\epsilon = 10$)	-6.88	*

Table 6: Relative error in the estimation of the delta value by each method. In experiment 1, (10) is used with $\gamma = 0.05$, and in experiment 2, (12) is used with $\gamma = 1$. “*” indicates that the method cannot work with the choice of distance function because it causes singularities in the computations.

variogram. In the previous experiments, we showed that we can have better results using a spherical variogram model; however, we haven’t provided any analysis supporting why this variogram is a better choice. In this section, we address this subject. In particular, we conduct experiments to explore whether we can increase the accuracy of the Kriging method by choosing a variogram function that can better approximate the empirical variogram.

To compute the empirical variogram, we partition the x-axis into 20 intervals of equal length $\frac{h_{\max}}{20}$ where h_{\max} is the maximum distance between two VA policies using the distance function (10) and with $\gamma = 1$. In each interval, to approximate (4), we use the average of the squared difference of the delta value of all pairs of VA policies that have a distance that falls into that interval as the representative for the empirical variogram for that interval. We call the piece-wise linear function that is formed by connecting the representative value for each interval the empirical variogram.

To approximate the empirical variogram, we use polynomials of degree 1, 2, 3 and 4. The polynomials are best MSE approximations of the empirical variogram in the interval between zero and range. At any distance above the range, the estimated value is assumed to be the value of the polynomial at the range (Figure 2). This is necessary in order to have a proper variogram model (Cressie, 1993) without producing jumps in the variogram.

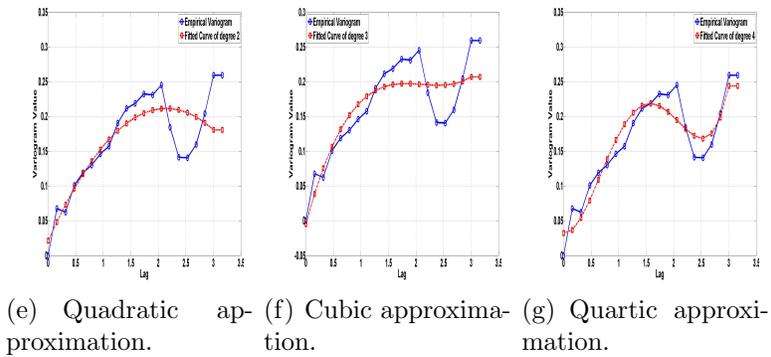
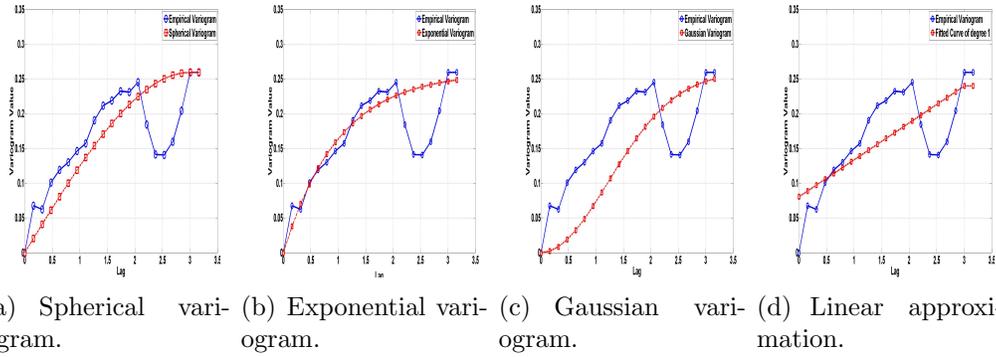


Figure 2: Comparing the variogram models with the empirical variogram.

The accuracy of Kriging using each approximation of the empirical variogram is presented in Table 7. An interesting, yet counter intuitive, observation is that the accuracy of Kriging is worst for the quartic MSE approximation variogram model that best fits the empirical variogram. Even comparing the graph of the exponential and spherical variogram models with the empirical variogram, the exponential variogram model seems to fit the empirical variogram model better than the spherical variogram model, but the accuracy of Kriging with the exponential variogram model is worse than the accuracy of Kriging with the spherical variogram model.

Because of these counter intuitive results, we took a closer look at the data from which the empirical variogram was generated. Figure 3 shows a graph of the squared differences of delta values of a pair of VA contracts versus their distance from each other. Surprisingly, the point values do not look similar to their average, i.e., the empirical variogram. We expected to see a graph similar to Figure 1 where the point values are in close proximity

Method	Relative Error (%)
Spherical	-0.03
Exponential	-1.61
Gaussian	< -500
Deg 1	-6.00
Deg 2	-32.17
Deg 3	-2.78
Deg 4	< -500

Table 7: Relative error in estimation of the delta value via Kriging with different variogram models.

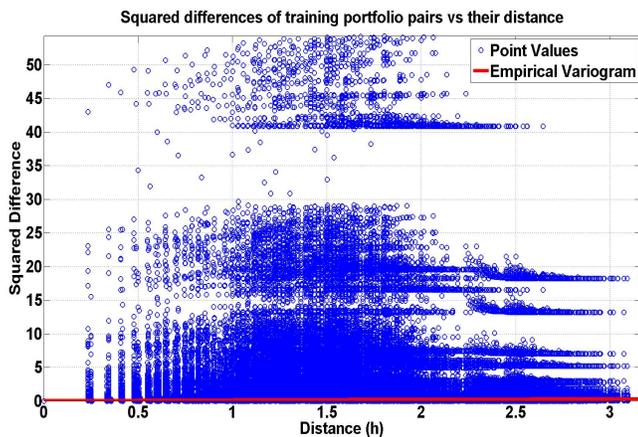


Figure 3: Squared difference of delta values of VA pairs in representative contracts.

to the empirical variogram and the variogram model. However, the data do not suggest the existence of any pattern from which a variogram model can be estimated. In particular, the data contradict the second-order stationary assumption underlying the Kriging method, and hence brings into question the appropriateness of the Kriging method for our application of interest.

5. Concluding Remarks

As we discussed in Sections 1 and 4, valuing a large portfolio of VA contracts via nested MC simulations is a time consuming process (Fox, 2013; Reynolds and Man, 2008). Recently, Gan and Lin (Gan, 2013a) proposed a Kriging framework that ameliorates the computational demands of the

valuation process. The proposed method is a special case of a more general framework called spatial interpolation. The spatial interpolation framework works as follows: A group of representative VA contracts is selected via a sampling method. For each contract in the sample, the Greeks are estimated via a MC simulation. Then a spatial interpolation method is used to find the Greeks for each contract in the portfolio as a linear combination of the Greeks of the contracts in the sample.

Two integral parts of the spatial interpolation framework are the choice of sampling method and the choice of interpolation scheme. In this paper, we studied various interpolation methods (i.e., Kriging, IDW, RBF) that can be used in this framework and postponed the discussion of the choice of sampling method to a future paper. In our experiments, we focused on a synthetic portfolio of VA contracts with GMDB and GMWB riders and compared the efficiency and accuracy of interpolation methods to estimate the delta value of this portfolio. As we discuss in detail in Section 4, while Kriging can provide better accuracy than the IDW and RBF methods, it has a slower running time than either of these deterministic methods, because Kriging has to solve (1) that in general takes $\Theta(n^3)$ time (n is the number of samples). The computational cost of Kriging is exacerbated if we compute the delta value of each VA policy in the portfolio. In general, a more granular view of the portfolio with Kriging comes at the expense of a much longer running time.

Our experiments also provided more insights into the importance of the choice of a distance function in the accuracy of various interpolation schemes. While Kriging methods, with exponential and spherical variogram models, can provide relatively accurate estimates using (10), they fail, due to numerical instability, to provide any estimates using (12). However, in comparison to (10), (12) allows IDW methods to provide better estimates. Our experiments also showed that achieving the best result with either (10) or (12) requires fine tuning the free parameters in these distance functions or in the methods themselves.

An interesting observation from our experiments is the non-existence of any pattern that supports the validity of the second-order stationary assumption that underlines the Kriging method. Our observation was the result of our experiments to accurately estimate the empirical variogram. Our hope was that a more accurate estimation of the empirical variogram would increase the accuracy of Kriging methods. Despite our initial belief, the experiments showed no relation between the accuracy of the Kriging method and

the closeness of the variogram model to the empirical variogram. Hence, it is not clear why the Kriging method provides accurate estimates despite the fact that the second-order stationary assumption does not seem to hold for the Greeks of the VA portfolio.

From our results, none of traditional spatial interpolation methods enjoy all three properties of accuracy, efficiency, and granularity. While Kriging based methods are accurate and fairly efficient at the portfolio level, they lack efficiency when a granular view of the portfolio is required. On the other hand, IDW and RBF based methods are efficient, and can provide a granular view of the portfolio, but they are not as accurate as Kriging based methods. In addition, the accuracy of all methods is dependent on the choice of an appropriate distance function, which requires tuning from expert users. It is not yet clear to us what is the best approach to find an appropriate distance function for each method. In a future paper, we will discuss how we can circumvent this issues and still achieve accuracy, efficiency, and granularity via a neural network approach to spatial interpolation. The proposed framework resolves the issues of the choice of distance function by learning an appropriate distance function to be used for each portfolio given the characteristics of the portfolio.

In the future, we also plan to address the problem of choosing an effective sampling method via a novel approach that uses statistical characteristics of the input portfolio to provide the output sample from the space in which the input portfolio is defined.

Acknowledgements

This research was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- Azimzadeh, P., Forsyth, P., 2013. The Existence of Optimal Bang-Bang Controls for GMxB Contracts. URL: https://cs.uwaterloo.ca/paforsyt/bang_bang.pdf.
- Bauer, D., Kling, A., Russ, J., 2008. A Universal Pricing Framework For Guaranteed Minimum Benefits in Variable Annuities. *ASTIN Bulletin* 38, 621–651.

- Bauer, D., Reuss, A., Singer, D., 2012. On the Calculation of the Solvency Capital Requirement Based on Nested Simulations. *ASTIN Bulletin* 42, 453–499.
- Bishop, C.M., 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, NJ, USA.
- Boyd, S., Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press, NY, USA.
- Boyle, P., Tian, W., 2008. The Design of Equity-Indexed Annuities. *Insurance: Mathematics and Economics* 43, 303–315.
- Boyle, P.P., Hardy, M.R., 1997. Reserving for Maturity Guarantees: Two Approaches. *Insurance: Mathematics and Economics* 21, 113–127.
- Burrough, P., McDonnell, R., Lloyd, C., 1998. *Principles of Geographical Information Systems*. 2nd ed., Oxford University Press.
- Belanger, A., Forsyth, P., Labahn, G., 2009. Valuing the Guaranteed Minimum Death Benefit Clause with Partial Withdrawals. *Applied Mathematical Finance* 16, 451–496.
- Carriere, J., 1996. Valuation of the Early-Exercise Price for Options Using Simulations and Nonparametric Regression. *Insurance: Mathematics and Economics* 19, 19–30.
- Cathcart, M., Morrison, S., 2009. Variable Annuity Economic Capital: the Least-Squares Monte Carlo Approach. *Life & Pensions* , 36–40.
- Chen, Z., Forsyth, P., 2008. A Numerical Scheme for the Impulse Control Formulation of Pricing Variable Annuities with a Guaranteed Minimum Withdrawal Benefit (GMWB). *Numerische Mathematik* 109, 535–569.
- Chen, Z., Vetzal, K., Forsyth, P., 2008. The Effect of Modelling Parameters on the Value of GMWB Guarantees. *Insurance: Mathematics and Economics* 43, 165–173.
- Chi, Y., Lin, X.S., 2012. Are Flexible Premium Variable Annuities Underpriced? *ASTIN Bulletin* 42, 559–574.

- Chiles, J.P., Delfiner, P., 1999. Geostatistics, Modelling Spatial Uncertainty. Wiley-Interscience.
- Coleman, T., L.Y., Patron, M.C., 2006. Hedging Guarantees in Variable Annuities under Both Equity and Interest Rate Risks. *Insurance: Mathematics and Economics* 38, 215–228.
- Cressie, N.A.C., 1993. *Statistics for Spatial Data*. John Wiley & Sons, Inc., New York, USA.
- Dai, M., Kwok, Y.K., Zong., J., 2008. Guaranteed Minimum Withdrawal Benefit in Variable Annuities. *Journal of Mathematical Finance* 18, 595–611.
- Daul, S., Vidal, E., 2009. Replication of Insurance Liabilities. *RiskMetrics Journal* 9, 79–96.
- Dembo, R., Rosen, D., 1999. The Practice of Portfolio Replication: A Practical Overview of Forward and Inverse Problems. *Annals of Operations Research* 85, 267–284.
- d’Halluin, Y., Forsyth, P., Vetzal, K., 2005. Robust Numerical Methods for Contingent Claims Under Jump Diffusion Processes. *IMA Journal of Numerical Analysis* 25, 65–92.
- Fox, J., 2013. A Nested Approach to Simulation VaR Using MoSes. *Insights: Financial modelling* , 1–7.
- Gan, G., 2013a. Application of Data Clustering and Machine Learning in Variable Annuity Valuation. *Insurance: Mathematics and Economics* 53, 795–801.
- Gan, G., 2013b. Representative Variable Annuity Policy Selection using Latin Hypercube Sampling. URL: http://bus.wisc.edu/~media/bus/knowledge-expertise/academic_departments/asrmi/workshops/spring-2014/gan2013latin.ashx.
- Gan, G., Lin, X.S., 2015. Valuation of Large Variable Annuity Portfolios Under Nested Simulation: A Functional Data Approach. *Insurance: Mathematics and Economics* 62, 138 – 150.

- Gan, G., Ma, C., Wu, J., 2007. *Data Clustering: Theory, Algorithms and Applications*. SIAM Press, Philadelphia, PA, USA.
- Gerber, H., Shiu, E., Yang, H., 2012. Valuing Equity-Linked Death Benefits and Other Contingent Options: A Discounted Density Approach. *Insurance: Mathematics and Economics* 51, 73–92.
- Hardy, M., 2003. *Investment Guarantees: Modeling and Risk Management for Equity-Linked Life Insurance*. John Wiley & Sons, Inc., Hoboken, New Jersey.
- Huang, Y., Forsyth, P., 2011. Analysis of A Penalty Method for Pricing a Guaranteed Minimum Withdrawal Benefit (GMWB). *IMA Journal of Numerical Analysis* 32, 320–351.
- Hull, J.C., 2006. *Options, Futures, and Other Derivatives*. 6th ed., Pearson Prentice Hall, Upper Saddle River, NJ.
- IRI, 2011. *The 2011 IRI fact book*. Insured Retirement Institute .
- Krige, D., 1951. *A Statistical Approach to Some Mine Valuations and Allied Problems at the Witwatersrand*. Master’s thesis, University of Witwatersrand.
- Lin, X., Tan, K., Yang, H., 2008. Pricing Annuity Guarantees under a Regime-Switching Model. *North American Actuarial Journal* 13, 316–338.
- Longstaff, F., Schwartz, E., 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. *The Review of Financial Studies* 14, 113–147.
- Matheron, G., 1963. Principles of Geostatistics. *Economic Geology* 58, 1246–1266.
- McKay, M., Beckman, R., Conover, W.J., 1979. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics* 21, 239–245.
- Milevsky, M., Salisbury, T., 2006. Financial Valuation of Guaranteed Minimum Withdrawal Benefits. *Insurance: Mathematics and Economics* 38, 21–38.

- Moenig, T., Baue, D., 2011. Revisiting the Risk-Neutral Approach to Optimal Policyholder Behavior: A study of Withdrawal Guarantees in Variable Annuities, in: 12th Symposium on Finance, Banking, and Insurance, Germany.
- Oechslin, J., Aubry, O., Aellig, M., Kappeli, A., Bronnimann, D., Tandonnet, A., Valois, G., 2007. Replicating Embedded Options in Life Insurance Policies. *Life & Pensions* , 47–52.
- Reynolds, C., Man, S., 2008. Nested Stochastic Pricing: The Time Has Come. *Product Matters!*— *Society of Actuaries* 71, 16–20.
- Shepard, D., 1968. A Two-dimensional Interpolation Function for Irregularly-spaced Data, in: *Proceedings of the 1968 23rd ACM National Conference*, ACM, NY, USA. pp. 517–524.
- Stein, E., Shakarchi, R., 2009. *Real Analysis: Measure Theory, Integration, and Hilbert Spaces*. Princeton University Press.
- TGA, 2013. Variable Annuities—An Analysis of Financial Stability. The Geneva Association
URL: <https://www.genevaassociation.org/media/ga2013/99582-variable-annuities.pdf>.
- Ulm, E., 2006. The Effect of the Real Option to Transfer on the Value Guaranteed Minimum Death Benefit. *The Journal of Risk and Insurance* 73, 43–69.