

Article

Use of Prediction Bias in Active Learning and Its Application to Large Variable Annuity Portfolios

Hyukjun Gweon ^{1,*}, Shu Li ¹ and Yangxuan Xu ²

¹ Department of Statistical and Actuarial Sciences, Western University, London, ON N6A 3K7, Canada; shu.li@uwo.ca

² Department of Statistical Science, Duke University, Durham, NC 27708, USA; yangxuan.xu@duke.edu

* Correspondence: hgweon@uwo.ca

Abstract: Given the computational challenges associated with valuing large variable annuity (VA) portfolios, a variety of data mining frameworks, including metamodeling and active learning, have been proposed in recent years. Active learning, a promising alternative to metamodeling, enhances the efficiency of VA portfolio assessments by adaptively improving a predictive regression model. This is achieved by augmenting data for model training with strategically selected informative samples. Successful application of active learning requires an effective metric in order to gauge the informativeness of data. Current sampling methods, which focus on prediction error-based informativeness, typically rely solely on prediction variance and assume an unbiased predictive model. In this paper, we address the fact that prediction bias can be nonnegligible in large VA portfolio valuation and investigate the impact of prediction bias in both the modeling and sampling stages of active learning. Our experimental results suggest that bias-based sampling can rival the efficacy of traditional ambiguity-based sampling, with its success contingent upon the extent of bias present in the predictive model.

Keywords: variable annuity; machine learning; active learning; prediction bias; data mining



Citation: Gweon, Hyukjun, Shu Li, and Yangxuan Xu. 2024. Use of Prediction Bias in Active Learning and Its Application to Large Variable Annuity Portfolios. *Risks* 12: 85. <https://doi.org/10.3390/risks12060085>

Academic Editor: Mogens Steffensen

Received: 1 April 2024

Revised: 2 May 2024

Accepted: 12 May 2024

Published: 22 May 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Variable annuities (VAs), also known as segregated fund contracts in Canada, are a popular type of equity-linked insurance product often utilized as investment vehicles within North American retirement savings plans. VAs typically offer a range of guaranteed benefits aimed at providing protection, such as the Guaranteed Minimum Death Benefit (GMDB), Guaranteed Minimum Maturity Benefit (GMMB), Guaranteed Minimum Accumulation Benefit (GMAB), and Guaranteed Minimum Withdrawal Benefit (GMWB) (see, e.g., Hardy (2003); Ledlie et al. (2008)). A substantial body of literature examines the stochastic modeling of embedded guarantees and their valuation in a fair market (risk-neutral) setting. However, the valuation of these embedded guarantees is complex due to the inherent complexity of diverse benefit designs. For example, GMDB guarantees a minimum amount to the policyholder upon death, exposing the insurer to various risks including investment, interest rate, mortality, and surrender risks. For a comprehensive overview of pricing, valuation, and risk management of individual VA contracts, we refer readers to existing literature, including Feng et al. (2022) and references therein. This paper primarily focuses on the valuation of large portfolios of VAs using statistical tools.

Due to the complex nature of the guarantees' payoffs, closed-form solutions are often not feasible. Consequently, insurance companies extensively utilize Monte Carlo (MC) simulations for valuation and dynamic hedging purposes. However, valuing a large portfolio of Variable Annuity (VA) contracts using Monte Carlo simulation is notably time-intensive, as it requires projecting each contract across numerous scenarios over an extended time period (Gan and Valdez 2018). To efficiently assess large VA portfolios

within constrained timeframes or resources, numerous data mining techniques have been introduced. The predominant framework, metamodeling, integrates both supervised and unsupervised machine learning methods through several sequential modeling stages. It employs unsupervised learning to generate a small, representative dataset, followed by supervised learning to fit a predictive model to these data after valuation is completed. Literature reviews reveal various unsupervised learning algorithms such as the truncated fuzzy c-means algorithm (Gan and Huang 2017), conditional Latin hypercube sampling (Gan and Valdez 2018; Minasny and McBratney 2006), and hierarchical k -means clustering (Gan and Valdez 2019). On the other hand, supervised learning approaches for VA portfolio valuation include kriging (Gan 2013), GB2 (Gan and Valdez 2018), group LASSO (Gan 2018), and several tree-based methods (Gweon and Li 2023; Gweon et al. 2020; Quan et al. 2021).

An alternative data mining approach is active learning (Gweon and Li 2021; Settles 2012), which involves iterative and adaptive data sampling and model training. Within this framework, the data sampling stage identifies a batch of informative policies, integrating these selected contracts into the existing representative data set. This augmented data set is then used for predictive model training. This cyclical process of sampling and updating continues until the allocated computational resources for portfolio assessment are exhausted. A prevalent strategy in active learning is uncertainty or ambiguity sampling (Burbidge et al. 2007; Settles and Craven 2008), which targets contracts where the current model shows significant predictive uncertainty.

Active learning methodologies have predominantly focused on classification problems, where the response variable is categorical. When the output variable is continuous, uncertainty can be quantified by prediction error, which can be decomposed into variance and squared bias. Current approaches to informativeness-based sampling for regression problems typically rely on prediction variance, under the assumption that the predictive model is approximately unbiased (Burbidge et al. 2007; Krogh and Vedelsby 1995; Kumar and Gupta 2020). However, as noted by Gweon et al. (2020), there are practical instances where prediction bias is significant and cannot be overlooked. This recognition that prediction bias is a critical component of prediction error has spurred interest in using prediction bias as a measure of uncertainty in active learning. This paper explores the potential of bias-based sampling techniques in light of these findings.

To the best of our knowledge, this study represents pioneering research that explores the application of bias-based sampling in active learning for regression problems. We introduce a metric for bias-based sampling and investigate its utility under conditions where the predictive model exhibits significant or minimal bias. Empirical evaluations on a large synthetic VA portfolio with six response variables demonstrate the effectiveness of bias-based sampling using random forest (RF, Breiman (2001)), particularly compared to variance-based sampling. However, the effectiveness diminishes when bias reduction techniques for RF are implemented.

The remainder of the paper is structured as follows: Section 2 reviews the active learning framework and discusses the estimation and application of prediction bias within this context. Section 2.4 applies the proposed methods to the efficient valuation of a large synthetic VA dataset. Section 3 concludes the paper.

2. Main Methodology

In this section, we propose several strategies to incorporate prediction bias into the active learning framework. We begin with a brief overview of the active learning framework and explore how prediction uncertainty is measured using random forest. Subsequently, we detail the integration of prediction bias into the modeling and sampling stages of active learning.

2.1. The Active Learning Framework for Large VA Valuation

Supervised statistical learning typically involves a one-time model training session using a complete set of labeled data. When large unlabeled data are available during the

model training phase, active learning that uses both labeled and unlabeled data offers a viable alternative to traditional supervised learning. The primary objective of active learning is to iteratively enhance the predictive model by incorporating additional, informative samples selected from the pool of unlabeled data. This approach is particularly advantageous for rapidly valuing large VA portfolios because there are a large number of unlabeled contracts available during the model training process. As described by [Gweon and Li \(2021\)](#), the active learning process comprises the following steps:

- (1) Formulate the initial representative data set by selecting a small subset from the entire portfolio, which can be chosen randomly.
- (2) Determine the values of the response variable for the selected contracts using Monte Carlo simulation, thereby generating the labeled data necessary for training the predictive model.
- (3) (Modeling stage) Train the predictive model using the labeled representative data.
- (4) (Sampling stage) Apply the predictive model to the remaining set of unlabeled VA contracts to identify a subset of contracts whose information could significantly improve the model.
- (5) Use a Monte Carlo simulation to label the newly selected contracts and integrate them into the existing set of labeled data, resulting in an expanded labeled dataset.
- (6) Repeat steps (3) through (5) until specific stopping criteria are met, typically when the allocated time for valuation expires.
- (7) Finalize the assessment by applying the updated predictive model to estimate the response values for the remaining unlabeled contracts.

The key part of an active learning algorithm is the selection of informative samples from the unlabeled data at each sampling iteration. Such informative samples lead to a great improvement in the current predictive model. Most existing approaches for informative sampling have been developed for classification problems and thus cannot be directly adapted to the VA valuation in which the target variables are continuous. Also, sampling a batch, instead of a single contract, is desirable for an efficient use of resources such as model runtime and parallel computing systems. Recently, [Gweon and Li \(2021\)](#) proposed batch-mode active learning approaches using random forest ([Breiman 2001](#)) for an efficient valuation of large VA portfolios. To quantify the informativeness of each VA contract, they adopted the principle of Query-By-Committee ([Abe and Mamitsuka 1998](#); [Burbidge et al. 2007](#); [Freund et al. 1997](#)) and used prediction ambiguity that is equivalent to the sample variance of bootstrap regression trees. They found that active learning can be highly effective when the informative samples are chosen by a weighted random sampling for which the sampling weights are proportional to prediction ambiguities.

A critical component of the active learning algorithm involves selecting informative samples from the unlabeled data during each iteration. These samples substantially enhance the accuracy and robustness of the existing predictive model. Traditional approaches for informative sampling, primarily developed for classification problems, are not directly applicable to VA valuation, where the target variables are continuous. Moreover, batch sampling—as opposed to selecting individual contracts—is desirable for an efficient use of resources, such as model runtime and parallel computing systems. Recently, [Gweon and Li \(2021\)](#) introduced batch-mode active learning strategies using random forest for the efficient valuation of large VA portfolios. To assess the informativeness of each contract, they implemented the Query-By-Committee principle ([Breiman 2001](#)) and measured prediction ambiguity, akin to the sample variance of bootstrap regression trees. Their findings indicate that active learning can be exceptionally effective when informative samples are selected through weighted random sampling, where the weights are proportional to the prediction ambiguities.

2.2. Random Forests and Quantifying Prediction Bias

Suppose that a VA portfolio contains a total of N contracts, $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ where each $\mathbf{x}_i \in \mathbb{R}^p$ represents a vector of feature variables associated with a contract. Random forest predicts the continuous response value Y_i for each contract \mathbf{x}_i using the regression form:

$$Y_i = f(\mathbf{x}_i) + \epsilon_i,$$

where $f(\cdot)$ represents the underlying model, and ϵ_i denotes random error. This modeling approach employs bootstrap regression trees (Breiman 1984) to estimate the regression function, with foundational concepts detailed in Gweon et al. (2020) and Quan et al. (2021).

Let L represent the labeled representative data containing n VA contracts. Random forest consisting of M regression trees constructs each regression tree $\hat{f}_{L_m}(\mathbf{x})$ using the m -th bootstrap sample L_m ($m = 1, \dots, M$). The predicted response value for an unlabeled contract \mathbf{x} is then derived by averaging across all M bootstrap regression trees:

$$\hat{f}(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M \hat{f}_{L_m}(\mathbf{x}).$$

The effectiveness of random forest in VA portfolio valuation has been demonstrated in prior studies (Gweon et al. 2020; Quan et al. 2021).

Given the necessity of quantifying prediction uncertainty in active learning, we focus on the mean square error (MSE), defined as:

$$\text{MSE}(\hat{f}(\mathbf{x})) = E\left((\hat{f}(\mathbf{x}) - f(\mathbf{x}))^2\right).$$

MSE is decomposed into the prediction (squared) bias and variance:

$$\begin{aligned} \text{MSE}(\hat{f}(\mathbf{x})) &= E\left((\hat{f}(\mathbf{x}) - E(\hat{f}(\mathbf{x})) + E(\hat{f}(\mathbf{x})) - f(\mathbf{x}))^2\right) \\ &= \left(E(\hat{f}(\mathbf{x})) - f(\mathbf{x})\right)^2 + E\left((\hat{f}(\mathbf{x}) - E(\hat{f}(\mathbf{x})))^2\right), \end{aligned} \quad (1)$$

where the first term represents the squared bias and the second term, the variance.

Although the averaging technique in random forest greatly improves the predictive performance of regression trees by reducing the prediction variance, it does not influence the prediction bias. With the response values in a VA portfolio potentially exhibiting a highly skewed distribution, Gweon et al. (2020) identified that the prediction bias of random forest is not negligible. Gweon et al. (2020) showed a bias-correction technique (Breiman 1999; Zhang and Lu 2012) can effectively estimate and correct prediction bias. Specifically, prediction bias can be assessed by fitting an additional random forest model to the out-of-bag (OOB) errors of the trained model. The OOB prediction for a data vector \mathbf{x}_i in the labeled set is defined as:

$$\hat{f}^{OOB}(\mathbf{x}_i) = \frac{1}{M_i} \sum_{m=1}^M \hat{f}_{L_m}(\mathbf{x}_i) I((\mathbf{x}_i, y_i) \notin L_m),$$

where $I(\cdot)$ denotes the indicator function, and M_i is the number of bootstrap samples not containing (\mathbf{x}_i, y_i) (i.e., $M_i = \sum_{m=1}^M I((\mathbf{x}_i, y_i) \notin L_m)$). A subsequent model $g(\cdot)$ is then trained where the response variable is the bias $B(\mathbf{x}) = \hat{f}^{OOB}(\mathbf{x}) - Y$, allowing the estimation of the prediction bias:

$$\hat{B}(\mathbf{x}) = \hat{g}(\mathbf{x}). \quad (2)$$

Combining the two random forest models, the bias-corrected random forest (BC-RF) predicts the response variable as the sum of $\hat{f}(\mathbf{x})$ and $\hat{g}(\mathbf{x})$. Gweon et al. (2020) utilized the quantity $\hat{B}(\mathbf{x})$ exclusively for mitigating prediction bias within the RF model as applied

in the metamodeling framework. In the following section, we discuss how this bias is leveraged within the active learning process.

2.3. Using Prediction Bias in Active Learning Process

Our focus is to explore the integration of prediction bias, $\hat{B}(\mathbf{x})$, into the sampling stage of the active learning process. Instead of developing a new sampling methodology, we adopt the weighted random sampling (WRS) approach from Gweon and Li (2021), which is already established as effective. Utilizing prediction bias within this framework is straightforward. Under WRS with prediction bias, the active learning algorithm selects a batch of unlabeled contracts of size d , where the sampling probability for a contract \mathbf{x}_i within the unlabeled data set U is defined as:

$$w_i = \frac{B_i^2}{\sum_h B_h^2},$$

where $B_i = \hat{B}(\mathbf{x}_i)$ for $i = 1, \dots, |U|$. We employ the squared bias terms to ensure non-negative sampling probabilities for all contracts and to facilitate a natural combination with the ambiguity component since the mean square error of the regression trees is the sum of variance and squared bias. We denote this approach as WRS-B to differentiate it from the ambiguity-based WRS (WRS-A), which calculates the weight as:

$$w_i = \frac{A_i}{\sum_h A_h},$$

where the ambiguity score $A_i = \frac{1}{M-1} \sum_{m=1}^M (\hat{f}_{L_m}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_i))^2$ represents the sample variance of the bootstrap regression trees, as detailed in Gweon and Li (2021).

When is WRS-B particularly effective in active learning? We posit that its effectiveness is contingent on the severity of prediction biases. If the predicted biases of the target response variable are significant, the weight w_i in WRS-B acts as a measure of relative prediction uncertainty for the i -th contract. Conversely, if predictions are nearly unbiased, the bias estimates (B_i) may merely introduce noise and offer no substantial information, rendering the use of these scores a less effective uncertainty measure.

This analysis suggests that WRS-B could serve as a promising uncertainty measure for RF in the valuation of large VA portfolios in the presence of prediction biases, as demonstrated by Gweon et al. (2020). A deeper analysis is required for BC-RF, which predicts the response variable using two sequential RF models. In the j -th iteration of active learning, the prediction by BC-RF is formulated as:

$$\hat{Y} = \hat{f}^j(\mathbf{x}) + \hat{g}^j(\mathbf{x}),$$

where $\hat{f}^j(\mathbf{x})$ is the response prediction by an RF model, and $\hat{g}^j(\mathbf{x})$ is the predicted bias by a subsequent RF model. If the primary contribution of WRS-B to the enhancement of $\hat{f}^{j+1}(\mathbf{x})$ in the $(j+1)$ -th iteration is bias reduction, the updated model $\hat{f}^{j+1}(\mathbf{x})$ will yield less biased predictions, potentially reducing the impact of $\hat{g}^{j+1}(\mathbf{x})$.

Furthermore, we need to examine the relative effectiveness of WRS-B compared to WRS-A. If prediction ambiguity and bias are highly correlated, both sampling approaches may yield similar outcomes. However, if they are not, the samples selected by WRS-B and WRS-A are likely to differ. Moreover, these two uncertainty components—ambiguity and bias—can be combined in various ways. Since their sum corresponds to the mean square error of the individual trees, another variant, WRS based on mean square error (WRS-MSE), uses the weight:

$$w_i = \frac{A_i + B_i^2}{\sum_h (A_h + B_h^2)}.$$

In Section 2.4, we further explore the application of prediction bias in active learning across various settings.

2.4. Application in Large Variable Annuity Portfolio Valuation

2.5. Description of Data

The data utilized in this study comprises a synthetic dataset from [Gan and Valdez \(2017\)](#), containing 190,000 variable annuity (VA) contracts. These contracts are characterized by 16 explanatory variables (14 continuous and 2 categorical), as summarized in Table 1. Each of the ten investment funds within the dataset is linked to one or more market indices, with the specific mappings to the five market indices detailed in Table 2.

Table 1. Summary statistics of the predictor attributes in the VA dataset.

Category	Description	Count		
Gender	M (Male Policyholder)	113,993		
	F (Female policyholder)	76,007		
Product type	ABRP (GMAB with return of premium)	10,000		
	ABRU (GMAB with annual roll-up)	10,000		
	ABSU (GMAB with annual ratchet)	10,000		
	DBAB (GMDB + GMAB with annual ratchet)	10,000		
	DBIB (GMDB + GMIB with annual ratchet)	10,000		
	DBMB (GMDB + GMMB with annual ratchet)	10,000		
	DBRP (GMDB with return of premium)	10,000		
	DBRU (GMDB with annual roll-up)	10,000		
	DBSU (GMDB with annual ratchet)	10,000		
	DBWB (GMDB + GMWB with annual ratchet)	10,000		
	IBRP (GMIB with return of premium)	10,000		
	IBRU (GMIB with annual roll-up)	10,000		
	IBSU (GMIB with annual ratchet)	10,000		
	MBRP (GMMB with return of premium)	10,000		
	MBRU (GMMB with annual roll-up)	10,000		
	MBSU (GMMB with annual ratchet)	10,000		
WBRP (GMWB with return of premium)	10,000			
WBRU (GMWB with annual roll-up)	10,000			
WBSU (GMWB with annual ratchet)	10,000			
Continuous	Description	Minimum	Mean	Maximum
gmwBalance	GMWB balance	0	35,611.54	499,708.73
gbAmt	Guaranteed benefit amount	0	326,834.59	1,105,731.57
FundValue1	Account value of the 1st investment fund	0	33,433.87	1,099,204.71
FundValue2	Account value of the 2nd investment fund	0	38,542.81	1,136,895.87
FundValue3	Account value of the 3rd investment fund	0	26,740.18	752,945.34
FundValue4	Account value of the 4th investment fund	0	26,141.80	610,579.68
FundValue5	Account value of the 5th investment fund	0	23,026.50	498,479.36
FundValue6	Account value of the 6th investment fund	0	35,575.67	1,091,155.87
FundValue7	Account value of the 7th investment fund	0	29,973.25	834,253.63
FundValue8	Account value of the 8th investment fund	0	30,212.11	725,744.64
FundValue9	Account value of the 9th investment fund	0	29,958.29	927,513.49
FundValue10	Account value of the 10th investment fund	0	29,862.24	785,978.60
age	Age of the policyholder	34.52	49.49	64.46
ttm	Time to maturity in years	0.59	14.54	28.52

For the response variables, we analyze six variables: the fair market values (FMV) and five deltas. The FMV of a VA contract is defined as the difference between the guaranteed payoff and the guaranteed cost. The delta measures the rate of change of the theoretical

option value relative to changes in the underlying asset’s price for a specific market index. The partial dollar deltas for a VA contract are calculated as follows:

$$\text{Delta}_i = \frac{FMV(PA_1, \dots, 1.01PA_i, \dots, PA_5)}{0.02} - \frac{FMV(PA_1, \dots, 0.99PA_i, \dots, PA_5)}{0.02},$$

where PA_i denotes the partial account value linked to the i -th market index and FMV is the fair market value as a function of partial account values. The partial dollar delta indicates the sensitivity of the guarantee value to fluctuations in the index, which is critical for determining appropriate hedge positions.

Figure 1 illustrates the distribution of the six response variables. Observations indicate that the distribution of FMV is significantly skewed to the right, while the distributions of the deltas are predominantly skewed to the left.

Table 2. Ten investment funds. Each row is a mapping from an investment fund to a combination of five indices.

Fund	US Large	US Small	Intl Equity	Fixed Income	Money Market
1	1	0	0	0	0
2	0	1	0	0	0
3	0	0	1	0	0
4	0	0	0	1	0
5	0	0	0	0	1
6	0.6	0.4	0	0	0
7	0.5	0	0.5	0	0
8	0.5	0	0	0.5	0
9	0	0.3	0.7	0	0
10	0.2	0.2	0.2	0.2	0.2

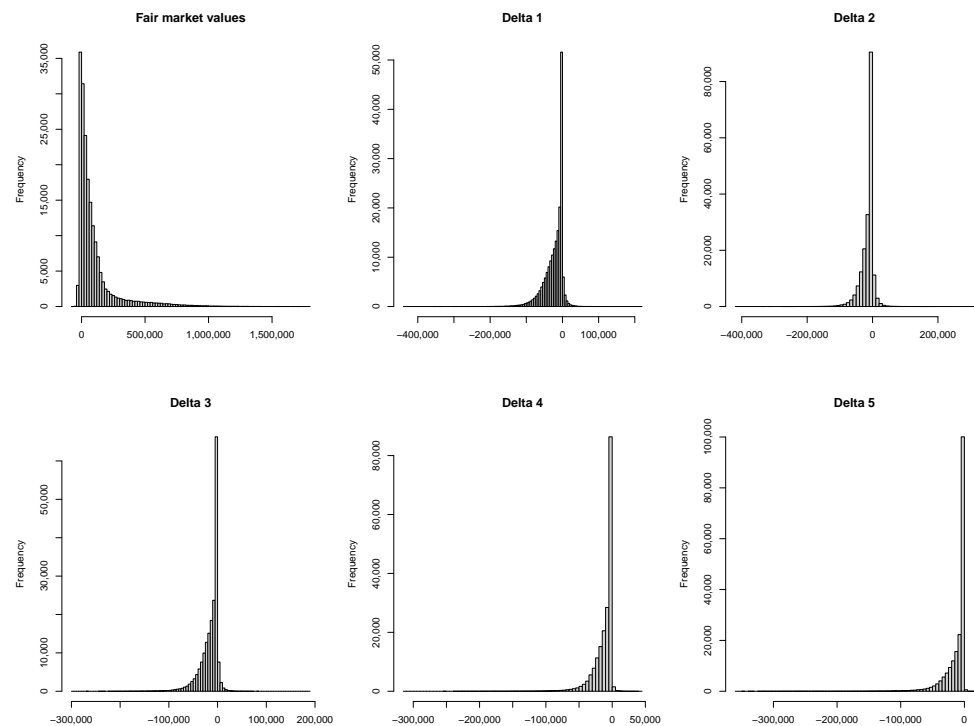


Figure 1. Distributions of the six response variables.

2.6. Experimental Setting

The active learning framework was initiated with a set of 100 representative contracts, selected through simple random sampling without replacement, to fit an initial regression model. For the response variable and the bias estimation, we employed two RF models, each using 300 regression trees to ensure stable performance, as supported by Gweon et al. (2020); Quan et al. (2021). Furthermore, consistent with findings from Quan et al. (2021), all 16 features were utilized in each binary split during tree construction, which has been shown to optimize performance for this dataset.

During each iteration of active learning, each sampling approach selected 50 contracts based on sampling probabilities. These contracts were then incorporated into the representative data set, and the RF models were subsequently updated with this augmented data. The iterative process was terminated once the representative data set expanded to include 800 contracts. To mitigate the impact of random variability, each active learning method was executed 20 times using different seeds. This experimental setup was consistently applied across all six response variables.

For each of the five partial deltas, it was noted that some contracts exhibited zero delta values, indicating no association with the relevant market index. Given that the link between a contract and its market index is predetermined, contracts with zero delta values were excluded from our prediction assessments to maintain the integrity and relevance of the analysis.

2.7. Evaluation Measures

To assess predictive performance, we employ several metrics: R^2 , mean absolute error (MAE), and absolute percentage error (APE):

$$R^2 = 1 - \frac{\sum_i (\hat{f}(\mathbf{x}_i) - f(\mathbf{x}_i))^2}{\sum_i (f(\mathbf{x}_i) - \bar{f}(\mathbf{x}))^2},$$

$$\text{MAE} = \frac{1}{N} \sum_i |\hat{f}(\mathbf{x}_i) - f(\mathbf{x}_i)|,$$

and

$$\text{APE} = \frac{|\sum_i f(\mathbf{x}_i) - \sum_i \hat{f}(\mathbf{x}_i)|}{\sum_i f(\mathbf{x}_i)},$$

where $\bar{f}(\mathbf{x}) = N^{-1} \sum_i f(\mathbf{x}_i)$. The R^2 and MAE measure accuracy at the individual contract level, while APE assesses portfolio-level accuracy, allowing for the offsetting of positive and negative prediction errors at the contract level. Higher values of R^2 indicate better performance, whereas lower values of MAE and APE are preferable.

2.8. Results

Initial investigations focused on identifying prediction bias across the six response variables. Figure 2 presents the out-of-bag (OOB) error of RF as a function of the true response variable. An observed increasing or decreasing pattern suggests a likely severe bias in RF predictions of the remaining data, underscoring the potential of prediction bias as a valuable uncertainty measure in active learning.

Figure 3 compares the performance of RF when using simple random sampling (RS), WRS-A and WRS-B. Both WRS-A and WRS-B consistently outperformed RS in terms of R^2 and MAE across all response variables as the active learning process progressed. However, in terms of APE, WRS-B tended to show more biased performance than RS. No clear advantage was observed between WRS-A and WRS-B; for instance, WRS-B yielded better R^2 results than WRS-A for FMV, while WRS-A excelled over WRS-B for delta 5. The two sampling methods performed comparably for other response variables.

The observed correlations between ambiguity and squared bias ranged from 0.4 to 0.5, indicating that WRS-A and WRS-B selected somewhat different contracts during

the sampling process. Figure 4 illustrates the average prediction standard errors and absolute biases under WRS-A and WRS-B, noting a gradual decrease across multiple iterations of active learning. Interestingly, WRS-A was more effective in reducing prediction standard errors than WRS-B, which excelled in diminishing prediction bias. These findings demonstrate distinct contributions of WRS-A and WRS-B to model improvement.

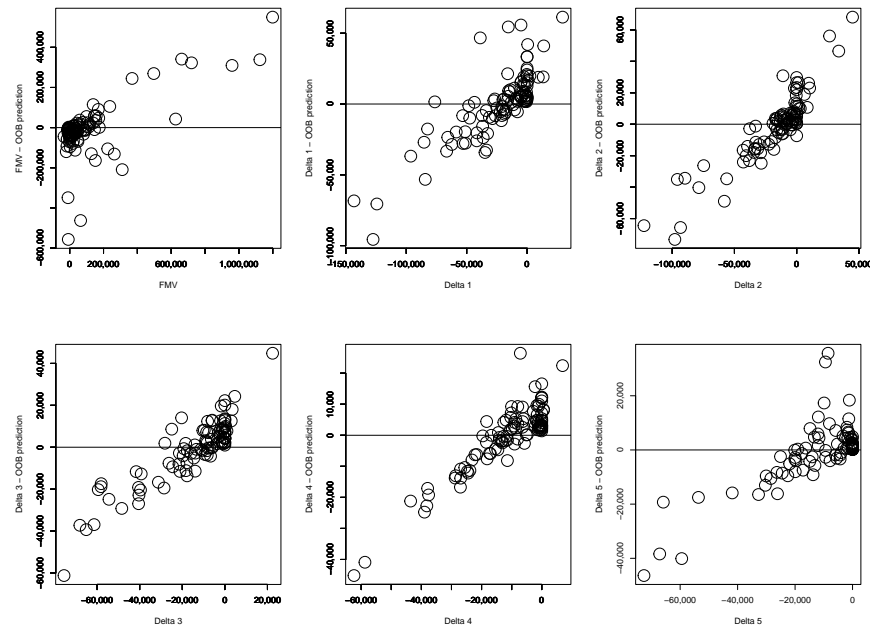


Figure 2. Out-of-bag prediction error results of RF on the initial representative data. Each panel corresponds to one of the six response variables.

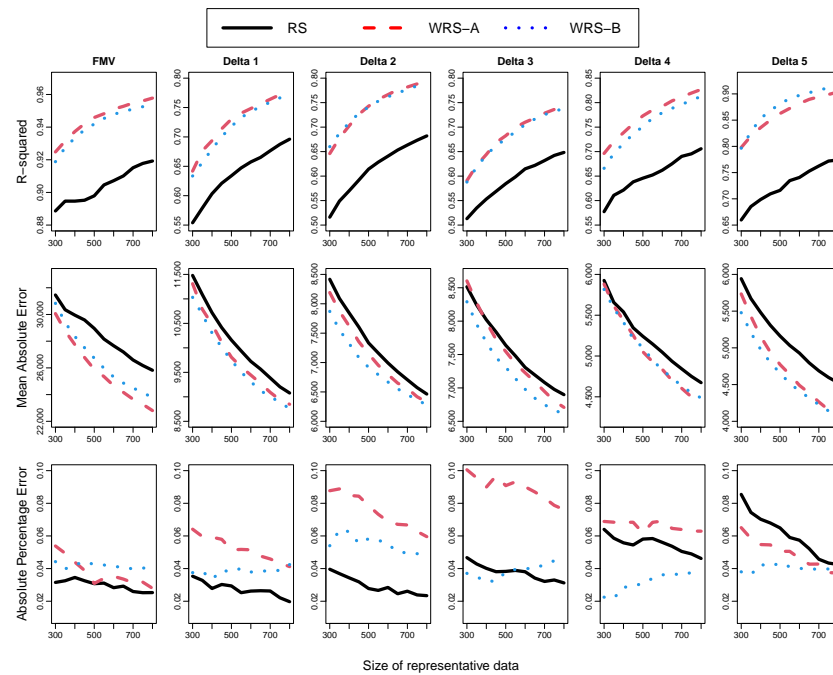


Figure 3. Predictive performance of RF under simple random sampling (solid), WRS-A (dotted), and WRS-B (dashed).

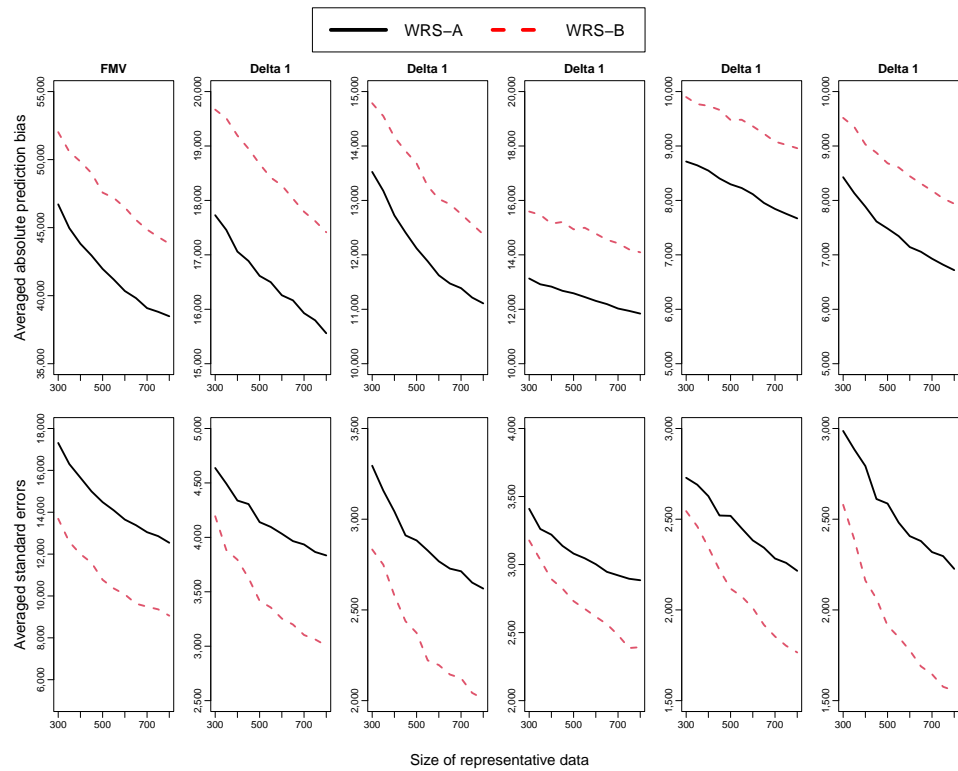


Figure 4. Averaged uncertainty of RF predictions for unlabeled data under WRS-A (dashed) and WRS-B (solid).

Subsequently, we explore the utilization of prediction bias in the modeling and sampling stages. Table 3 compares the performance of RF with BC-RF under each sampling approach. The introduction of bias correction in the modeling stage significantly enhanced predictive performance across all metrics. Notably, the large APE previously observed in RF under both WRS-A and WRS-B were considerably reduced in BC-RF. Given these comparative results, BC-RF is recommended for use in active learning regardless of the sampling strategy.

We further examine the performance of BC-RF using WRS-A, WRS-B, and WRS-MSE in Figure 5. Contrary to the inconclusive results observed with RF where no distinct advantage was noted between WRS-A and WRS-B, the implementation of BC-RF in the modeling stage generally showed WRS-A to be more effective than WRS-B in most situations. This observation aligns with earlier findings indicating that WRS-B substantially contributes to bias reduction in RF (refer to Figure 4). Given that RF under WRS-B is relatively less biased compared to its counterpart under WRS-A, the subsequent application of BC-RF yields smaller gains from further bias reduction. This outcome is evidenced in Tables 4 and 5, where the transition from RF to BC-RF exhibits more pronounced improvements under WRS-A.

Table 6 presents the runtime for each method until the representative VA contracts reached 800. When utilizing RF, WRS-A required approximately half the time compared to the other methods, as it did not involve an additional RF model for estimating prediction bias. In contrast, when BC-RF was employed, all sampling methods showed similar runtimes.

Interestingly, WRS-MSE often matched or exceeded the performance of WRS-A, as noted in the cases of Delta 1 and Delta 4. Considering that the computational times for WRS-MSE and WRS-A are nearly identical, these findings support a preference for WRS-MSE. Based on our analysis, we recommend employing BC-RF in the modeling stage alongside WRS-A or WRS-MSE in the sampling stage as an effective active learning strategy.

Table 3. Predictive performance of RF and BC-RF under each sampling approach ($N = 800$).

Response	Sampling	R^2		MAE		APE	
		RF	BC-RF	RF	BC-RF	RF	BC-RF
FMV	RS	0.9192	0.9480	25821.807	21142.456	0.0253	0.0176
	WRS-A	0.9537	0.9690	23813.411	19586.261	0.0409	0.0104
	WRS-B	0.9578	0.9650	22804.387	20953.240	0.0279	0.0119
Delta 1	RS	0.6962	0.7396	9078.134	7812.380	0.0197	0.0143
	WRS-A	0.7729	0.8114	8761.108	7469.906	0.0425	0.0139
	WRS-B	0.7761	0.8062	8847.688	7877.713	0.0412	0.0135
Delta 2	RS	0.6821	0.6735	6464.244	5831.808	0.0234	0.0171
	WRS-A	0.7903	0.8278	6282.434	5448.818	0.0493	0.0194
	WRS-B	0.7949	0.8173	6326.927	5742.336	0.0596	0.0185
Delta 3	RS	0.6483	0.6926	6899.186	5992.517	0.0312	0.0175
	WRS-A	0.7388	0.7928	6603.941	5544.816	0.0490	0.0132
	WRS-B	0.7443	0.7802	6708.572	5925.164	0.0762	0.0310
Delta 4	RS	0.7059	0.7470	4673.581	3973.629	0.0462	0.0194
	WRS-A	0.8126	0.8647	4486.339	3573.041	0.0395	0.0124
	WRS-B	0.8270	0.8574	4413.975	3905.542	0.0629	0.0304
Delta 5	RS	0.7742	0.7766	4524.483	3960.967	0.0425	0.0159
	WRS-A	0.9145	0.9311	4072.907	3438.134	0.0395	0.0079
	WRS-B	0.9038	0.9111	4091.317	3716.284	0.0373	0.0142

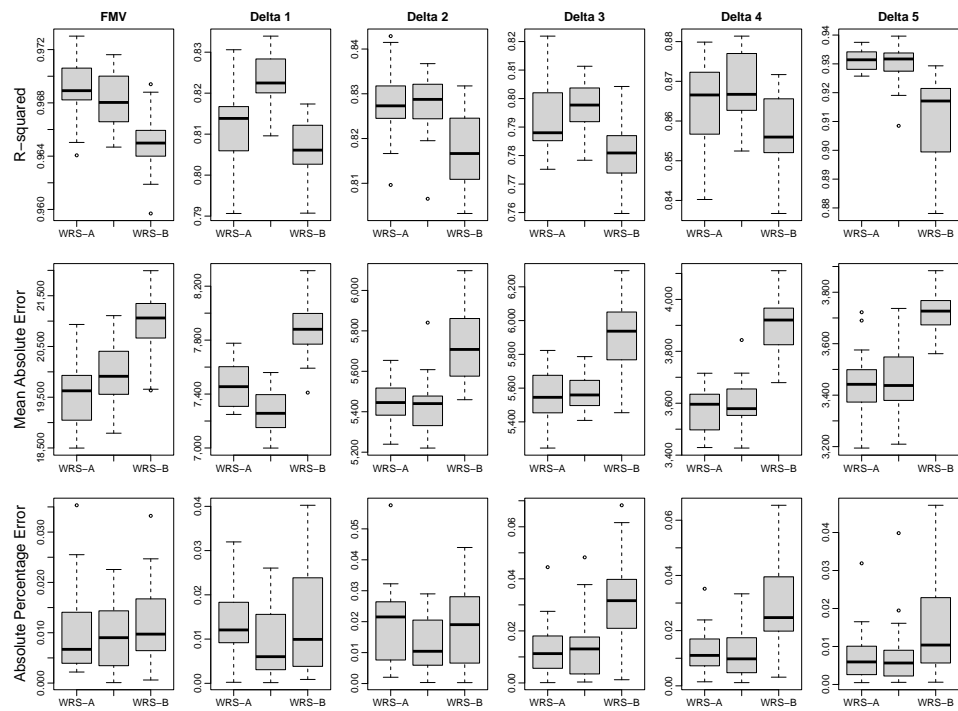


Figure 5. Boxplot results of BC-RF under different active learning sampling approaches.

Table 4. Relative improvement (in %) obtained by switching RF to BC-RF under WRS-A and WRS-B ($N = 500$).

	R^2		MAE		APE	
	WRS-A	WRS-B	WRS-A	WRS-B	WRS-A	WRS-B
FMV	1.80	0.87	15.54	8.96	60.54	32.84
Delta 1	7.46	3.15	15.42	9.18	46.31	60.88
Delta 2	5.18	3.63	11.88	9.29	54.47	67.90
Delta 3	8.98	6.50	14.86	11.42	47.05	54.50
Delta 4	8.04	4.37	18.32	10.08	44.98	41.40
Delta 5	2.07	1.49	15.22	9.60	63.55	62.05

Table 5. Relative improvement (in %) obtained by switching RF to BC-RF under WRS-A and WRS-B ($N = 800$).

	R^2		MAE		APE	
	WRS-A	WRS-B	WRS-A	WRS-B	WRS-A	WRS-B
FMV	1.60	0.75	17.75	8.12	74.47	57.29
Delta 1	4.99	3.88	14.74	10.96	67.19	67.15
Delta 2	4.75	2.82	13.27	9.24	60.65	68.93
Delta 3	7.31	4.84	16.04	11.68	72.99	59.30
Delta 4	6.41	3.68	20.36	11.52	68.48	51.72
Delta 5	1.82	0.81	15.59	9.17	79.96	61.90

Table 6. Runtime (in seconds) of each active learning approach. Each active learning process was terminated at $N = 800$.

	RF			BC-RF		
	WRS-A	WRS-B	WRS-MSE	WRS-A	WRS-B	WRS-MSE
FMV	61.63	128.38	128.37	130.74	128.52	130.94
Delta 1	64.00	130.33	133.53	133.32	135.07	132.94
Delta 2	57.48	119.00	122.13	122.50	120.47	119.65
Delta 3	59.28	124.65	125.05	119.17	122.96	122.27
Delta 4	52.71	107.24	108.37	109.23	109.72	109.74
Delta 5	52.82	101.41	102.77	104.10	103.87	104.00

3. Concluding Remarks

In this study, we explored the integration of prediction bias within the active learning framework, specifically for the valuation of large variable annuity portfolios. Using a synthetic portfolio with six distinct response variables, we assessed the utility of prediction bias in both the modeling and sampling stages of active learning. Our experimental findings indicate that (a) employing prediction bias for uncertainty sampling (WRS-B) enhances performance of RF in terms of R^2 and MAE, proving more effective than random sampling and as competitive as ambiguity-based uncertainty sampling (WRS-A); (b) uncertainty sampling based on prediction bias significantly contributes to the reduction of bias in RF; (c) prediction bias proves beneficial in the modeling stage across different sampling methods; and (d) WRS-B tends to be less effective than WRS-A when bias reduction is actively pursued during the modeling phase.

The batch-mode active learning framework necessitates an initial set of labeled representative data and a defined batch size. Although we employed simple random sampling to generate the initial datasets, alternative strategies involving unsupervised learning algorithms might be viable. Our tests with conditional Latin hypercube sampling indicated that the unsupervised learning method did not significantly influence the outcomes within the active learning framework. Furthermore, the batch size inversely affects the runtime of the active learning process; a larger batch size decreases the number of iterations needed to achieve the desired dataset size, as more data are labeled and incorporated at each

step. Our analysis, initially based on a batch size of 50, showed that increasing this to 100 maintained similar performance across different sampling methods.

In conclusion, this paper highlights the potential of prediction bias as an effective uncertainty measure within active learning for actuarial applications. While our approach utilized RF, other statistical learning algorithms could be considered if prediction bias can be appropriately quantified. A comprehensive examination of bias-based active learning employing various statistical methods merits further investigation.

Author Contributions: Conceptualization, H.G. and S.L.; methodology, H.G.; software, Y.X.; validation, H.G., S.L. and Y.X.; formal analysis, S.L. and Y.X.; investigation, Y.X.; resources, H.G.; data curation, Y.X.; writing—original draft preparation, H.G. and Y.X.; writing—review and editing, H.G. and S.L.; visualization, Y.X.; supervision, H.G. and S.L.; project administration, H.G. and S.L.; funding acquisition, H.G. and S.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC, grant numbers 06219 (PI: Shu Li) and 04698 (PI: Hyukjun Gweon)).

Data Availability Statement: The synthetic dataset used in this article is available at <https://www2.math.uconn.edu/~gan/software.html> (accessed on 12 May 2024).

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Abe, Naoki, and Hiroshi Mamitsuka. 1998. Query learning strategies using boosting and bagging. Paper presented at the 15th International Conference on Machine Learning (ICML), Madison, WI, USA, July 24–27; pp. 1–9.
- Breiman, Leo. 1984. *Classification and Regression Trees*. Boca Raton: Taylor & Francis, LLC.
- Breiman, Leo. 1999. *Using Adaptive Bagging to Debias Regressions*. Technical Report; Berkeley: University of California at Berkeley.
- Breiman, Leo. 2001. Random forests. *Machine Learning* 45: 5–32. [CrossRef]
- Burbidge, Robert, Jem J. Rowland, and Ross D. King. 2007. Active learning for regression based on query by committee. In *Intelligent Data Engineering and Automated Learning-IDEAL 2007*. Edited by H. Yin, P. Tino, E. Corchado, W. Byrne and X. Yao. Berlin and Heidelberg: Springer, pp. 209–18.
- Feng, Runhuan, Guojun Gan, and Ning Zhang. 2022. Variable annuity pricing, valuation, and risk management: a survey. *Scandinavian Actuarial Journal* 2022: 1–34. [CrossRef]
- Freund, Yoav, H. Sebastian Seung, Eli Shamir, and Naftali Tishby. 1997. Selective sampling using the query by committee algorithm. *Machine Learning* 28: 133–68. [CrossRef]
- Gan, Guojun. 2013. Application of data clustering and machine learning in variable annuity valuation. *Insurance: Mathematics and Economics* 53: 795–801.
- Gan, Guojun. 2018. Valuation of large variable annuity portfolios using linear models with interactions. *Risks* 6: 71. [CrossRef]
- Gan, Guojun, and Jimmy Xiangji Huang. 2017. A data mining framework for valuing large portfolios of variable annuities. Paper presented at the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13–17; pp. 1467–75.
- Gan, Guojun, and Emiliano A. Valdez. 2017. Valuation of large variable annuity portfolios: Monte carlo simulation and synthetic datasets. *Dependence Modeling* 5: 354–74. [CrossRef]
- Gan, Guojun, and Emiliano A. Valdez. 2018. Regression modeling for the valuation of large variable annuity portfolios. *North American Actuarial Journal* 22: 40–54. [CrossRef]
- Gan, Guojun, and Emiliano A. Valdez. 2019. Data clustering with actuarial applications. *North American Actuarial Journal* 1–19. [CrossRef]
- Gweon, Hyukjun, and Shu Li. 2021. Batch mode active learning framework and its application on valuing large variable annuity portfolios. *Insurance: Mathematics and Economics* 99: 105–15. [CrossRef]
- Gweon, Hyukjun, and Shu Li. 2023. A hybrid data mining framework for variable annuity portfolio valuation. *ASTIN Bulletin: The Journal of the IAA* 53: 580–95.
- Gweon, Hyukjun, Shu Li, and Rogemar Mamon. 2020. An effective bias-corrected bagging method for the valuation of large variable annuity portfolios. *ASTIN Bulletin: The Journal of the IAA* 50: 853–71. [CrossRef]
- Hardy, Mary. 2003. *Investment Guarantees: Modelling and Risk Management for Equity-Linked Life Insurance*. Hoboken: John Wiley & Sons, Inc.
- Krogh, Anders, and Jesper Vedelsby. 1995. Neural network ensembles, cross validation, and active learning. In *Advances in Neural Information Processing Systems*. Cambridge: MIT Press, pp. 231–38.
- Kumar, Punit, and Atul Gupta. 2020. Active learning query strategies for classification, regression, and clustering: A survey. *Journal of Computer Science and Technology* 35: 913–45. [CrossRef]

- Ledlie, M. Colin, Dermot P. Corry, Gary S. Finkelstein, Alan J. Ritchie, Ken Su, and D. C. E. Wilson. 2008. Variable annuities. *British Actuarial Journal* 14: 327–89. [[CrossRef](#)]
- Minasny, Budiman, and Alex B. McBratney. 2006. A conditioned latin hypercube method for sampling in the presence of ancillary information. *Computers & Geosciences* 32: 1378–88.
- Quan, Zhiyu, Guojun Gan, and Emiliano Valdez. 2021. Tree-based models for variable annuity valuation: parameter tuning and empirical analysis. *Annals of Actuarial Science* 16: 95–118. [[CrossRef](#)]
- Settles, Burr. 2012. *Active Learning*. San Rafael: Morgan & Claypool.
- Settles, Burr, and Mark Craven. 2008. An analysis of active learning strategies for sequence labeling tasks. Paper presented at the Conference on Empirical Methods in Natural Language Processing, Honolulu, HI, USA, October 25–27; pp. 1070–79.
- Zhang, Guoyi, and Yan Lu. 2012. Bias-corrected random forests in regression. *Journal of Applied Statistics* 39: 151–60. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.