THE PENNSYLVANIA STATE UNIVERSITY
SCHREYER HONORS COLLEGE

DEPARTMENT OF RISK MANAGEMENT

MARGINAL PRICING:
AN EDUCATIONAL MODEL TO PRICE CORRELATED RISKS IN INSURANCE

MICHAEL G. CALLAHAN
SPRING 2019

A thesis
submitted in partial fulfillment
of the requirements
for a baccalaureate
in MATHEMATICS WITH ACTUARIAL OPTION
with honors in MATHEMATICS

Reviewed and approved* by the following:

NAN ZHU
ASSISTANT PROFESSOR OF RISK MANAGEMENT
THESIS SUPERVISOR

LUEN-CHAU LI
PROFESSOR OF MATHEMATICS
HONORS ADVISOR

DENNIS PEARL
RESEARCH PROFESSOR OF STATISTICS
FACULTY READER

* Signatures are on file in the Schreyer Honors College.

ABSTRACT

This paper expands upon existing educational actuarial models to construct a new general insurance pricing model that will demonstrate the effects of additional actuarial concepts to students via an interactive web application built in R-Shiny. The new model is designed to calculate the marginal cost of adding one additional policy, or group of policies, to an existing business portfolio of policies. The paper achieves this by viewing all policies from an investment portfolio perspective and projecting the resulting cashflows forward in time using simulation. Then, the model assesses the simulation and calculates the average net present value (NPV) of the simulation, which is the sole financial objective for the company represented by the model. The paper defines the "fair upfront marginal premium" of a given policy to be the change in average NPV caused by adding the additional policy. Thus, the premium can be considered the upfront premium rate at which the insurer's utility is the same with the additional insured as without. The resulting model is a simplification of the different actuarial measurements taken to effectively price insurance in industry, but serves as an apt educational framework for students to understand more practical business applications. Particularly, the key advantage of the model is in revealing deeper patterns with respect to the cost of policy loss variance. Ultimately, this paper finds that under such a pricing framework, the risk appetite of an insurance company varies dramatically depending upon its current financial circumstance, and that this model can provide revealing metrics of implied risk premiums to provide objective measures of risk appetite.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

I would like to thank Nan Zhu for supporting my thesis ideas, encouraging me to create something new, and appreciating the value in an academic exercise. Additionally, the honors option reading assignment Dr. Zhu gave me on Mean-Variance portfolio theory serves as the inspirational cornerstone of my thesis. Most of the tools used in this paper are the ones he gave me.

I would also like to thank Dennis Pearl, Ron Gebhardstbauer, Matthew Katz, and Nestor Handzy for being excellent professors and helping me grow as a mathematics student and an actuary. All the time and effort they put into welcoming and answering so many of my questions during office hours have made me a more intellectually curious person.

Additionally, I would like to thank my peers in RM 494H, who have helped me reshape and revise my thesis, even when my sentences were 10 lines long.

Finally, I would like to thank my family for all their support and encouragement during my time at Penn State.

Chapter 1

Introduction

Insurers are incentivized to compete on the price of policies written. The actuarial practice of pricing policies is a complex process influenced by simplifying assumptions, including those made about the following: the nature of policies to be sold; the costs of policies already sold; policyholder behavior; and investments made over the term of the policy. As a result, the modeling techniques employed to price policies in industry are likewise complex and varied. However, by the time most undergraduate actuarial students find professional employment, they have little exposure to the application of these various techniques which support the basis of the actuarial profession. To increase undergraduate exposure to these peripheral pricing techniques, an educational pricing model is created and built into an accessible tool for students to explore as a supplementary academic exercise. In particular, model is built to provide deeper insights into the risk appetites of insurance firms.

In Chapter 2, the literature review section of the paper achieves two objectives. The first objective is to motivate students to contemplate the various business needs that actuarial models must service in a real-world setting. The second objective is to review the academic background provided by undergraduate institutions following the standard Society of Actuaries (SOA) curriculum, with a focus on the two primary pricing techniques covered in the Long Term Actuarial Mathematics (LTAM) curriculum; the Equivalence Principle, and the Portfolio Percentile Principle. While mastery of these two techniques is indeed necessary to understanding real-world insurance pricing, it is insufficient to grasp the wider scope of actuarial techniques

that are used in industry to support such pricing decisions. As a result, this section will also provide a secondary focus on the supporting techniques of stochastic Monte Carlo simulation and portfolio theory covered in the new Investment and Financial Markets (IFM) exam curriculum. The model created in this paper seeks to borrow theoretical concepts from these more varied techniques and combine them with the two simpler premium calculation methods to price basic general insurance policies. The result is a novel pricing model which incorporates the additional considerations of solvency, investment risk and loss interdependency under a single framework.

In Chapter 3, the paper discusses the new educational model in detail. First, the paper posits a simple hypothetical example featuring an existing insurance company to motivate the analysis. Simplifying assumptions regarding the company's investment returns, existing assets and liabilities and financial objectives are first made both to clarify the scope of the problem and to provide quantitative inputs for variation. Once these assumptions are made, the paper proposes a stochastic model as a solution by outlining its methodology. Namely, discrete time stochastic simulations are run in which net cash flows from the existing base of assets and liabilities are first forecasted into the future to project future company equity. These values are then discounted back to the present day and added together to form a path-dependent Net Present Value (NPV) calculation. Statistical measurements are taken from the distribution of simulated NPV calculations. Then, the existing portfolio base of assets and liabilities is modified to accommodate the additional policyholders being priced by the model, and the simulation is run again. The difference in mean simulated NPVs from the *a priori*, or original, portfolio base to the *a posteriori*, or adjusted portfolio base, serves as the net upfront marginal premium of the model. Once the model dynamics are discussed, mathematical analyses are conducted to better

understand the behavior of the model, including recursive definitions of the probability density function and estimates of variance-convergence. Lastly, the model is coded into an interactive R-Shiny web application to serve as educational tool, and the chapter ends by discussing the utility of these applications.

In Chapter 4, the quantitative results are pulled from the web application to analyze the deeper relationships of the model. Specifically, univariate relationships are graphed to expose the effects of isolated assumption shocks on policy price. Patterns are drawn from this analysis and extended to provide insights into general insurance business strategy. After this, the paper will turn its focus to the secondary purpose of this model; to explore the relationship between risk appetite and annual loss ratio.

Finally, in Chapter 5, the model and its associated tools are discussed to expose their greatest strengths and weaknesses, and conclusions are drawn regarding the model's practical applicability and educational value. These insights will be followed by suggestions for further investigations into the model.

Chapter 2

Literature Review

The literature review of this paper serves two purposes: first, to outline the practical business needs that an actuarial pricing model should service; and second, to review all prerequisite undergraduate actuarial mathematics for understanding the model.

2.1 Pricing Motivation

Selling a product in an open market – be it life insurance or a steak knife – might seem quite simple. One must first consider the cost of producing the good sold, and then add a healthy padding of profit margin while still trying to keeping the price competitive enough to sell a high volume of that good. However, we can further separate this business problem into two distinct challenges – estimating the cost of goods sold, or the supply curve, and estimating the consumer demand curve. While both forces drive market prices equally through the law of supply and demand, it is often the sole focus of academic actuarial models accomplish the former. Thus, the rest of the paper will reduce the scope of the problem of pricing insurance to merely estimating the cost of goods sold. Within this scope, however, insurance is still a relatively complex service to price. There are three primary needs that any practical insurance pricing system must serve: the need for a holistic reflection of the annual loss risk profile; the need for future business solvency; and the need for pricing precision.

To capture the effect on annual loss risk profile that a given policy may have, one must take precise measurements of the loss distribution of policies annually, including the annual loss' mean and variance. While many practical pricing models may claim that these measurements can be safely taken

under the assumption of policy loss independence, this is not the case for numerous types of insurance policies. Cyber Risk insurance, for example, may cover the damages caused by the hacking of various companies which all share the same cyber security infrastructure. Should one such company be breached, the likelihood of the others in the portfolio also being breached would be raised significantly, implying a loss correlation. Loss interdependency can even be argued in insurance markets like health and life insurance, where mortality and morbidity improvements are made idiosyncratically rather than systematically. For example, the life expectancies of two different diabetics may hinge on the same development of diabetes treatments, making them more closely related in risk than someone who is not diabetic. However, due to the larger number of policies in these markets, the ultimate total loss each year, given by the law of large numbers, will be largely unaffected by such relationships. Thus, we can see that any model capturing loss interdependency can improve the estimation of annual loss distributions for some, but not all, insurance portfolios.

The next concern for a pricing model is future business solvency. Simply put, the policies must be priced so that the insurer can expected sustained earnings. Should the losses in any future year of the policy be greater than the premiums earned, a reserve must be used to pay off the difference. Should the company not predict this future loss and not set aside a large enough reserve, it might be forced to file for bankruptcy due to insolvency. Obviously, the goal of any pricing model would be to expose and study this risk to prevent such an occurrence. And although there are regulations put in place which require such adequate reserves to be held, an accurate risk assessment is nonetheless required for proper financial planning.

Finally, there is something to be said about the motivation for model precision. One might ask "why is it important to differentiate between policies and charge different prices for them, when it is far easier to assign all policies the same price?" The simple answer is to guard against adverse selection. Adverse selection "occurs when an insurer's premium revenues are insufficient to cover its insured losses because it insures an increasing proportion of high-risk customers and a dwindling proportion of low risk

customers" (Cather, 2018). If you charge all your insureds the same price, but at least one other competing company in the market differentiates on price between high and low risk customers, then logically your low risk customers will wish to find a better deal on insurance from your competitor, while the high-risk insureds in the market will wish to find a better deal on insurance from you. To prevent this, insurance companies must compete on price, and more specifically, the precision of their pricing models.

With these three practical business needs, we may formulate criteria for a useful educational pricing model for undergraduate actuarial science students. Keep the above needs in mind when reading the next section of the paper, which covers pre-existing educational models, to evaluate their strengths and weaknesses.

2.2 Academic Background

This section will review the academic knowledge required to understand the educational model of the paper. Knowledge of elementary calculus (MATH 230), probability theory (MATH 414) and financial mathematics (RM 410) will be required to begin the review. The pricing methods used in the SOA's LTAM examination will be covered first, using the textbook *Actuarial Mathematics for Life Contingent Risks* as the definite source of information. Then, Monte Carlo valuation methods will be introduced using the SOA's Investment and Financial Markets (IFM) exam official text, *Derivatives Markets*, as the primary source of information. Finally, Mean-Variance Portfolio Theory will be summarized from the SOA's official supplementary IFM text, *Corporate Finance*. This will complete the conceptual review for undergraduate students endeavoring to study the model.

2.2.1 Long-Term Actuarial Mathematics

To begin a review of various methods of premium calculation, we must first define some key terms. According to the LTAM text, "an **insurance policy** is a financial agreement between the insurance company and the policyholder [also called an 'insured']. The insurance company agrees to pay some benefits, for example the sum insured on the death of the policyholder within the term of an insurance policy, and the policyholder agrees to pay premiums to the insurance company to secure these benefits" (Dickson, 2013, p. 144). The purpose of the premium is preeminently to help the insurance company recoup the losses associated with said policy. A premium does not necessarily have to reflect the cost of additional policy expenses, such as sales commission, underwriting expenses, or business overhead. A premium calculated to reflect only the cost of policy benefits is called a **"net premium"** (Dickson, 2013)**.** A premium calculated to reflect costs of both policy benefits and additional expenses is called a **"gross premium"** (Dickson, 2013)**.** Additionally, premiums may be calculated on the basis of a one-time payment, called a **"single premium,"** or on the basis of regular installments (Dickson, 2013). This paper focuses on the **"net single premium"** calculation as the basis for generating a new pricing model (Dickson, 2013).

For a given policy, we denote L_0^n as "**net future loss"** (Dickson, 2013, p. 145)**.** This is equal to the present value of all future cash outflows due to benefits, less the present value of all premiums collected. For example, suppose an insured purchases a one-year $15,000 death benefit life insurance policy for $200 in a single payment. Assume the payment will come at the end of the year if the insured dies. With an interest rate of 5%, let's calculate L_0^n if (a) the insured does not die within the year, and (b) if the insured does die:

(a) $L_0^n = PV(benefits) - PV(premiums) = 0 - 200 = -200$

(b) $L_0^n = PV(benefits) - PV(premiums) = 15{,}000/(1.05) - 200 = 14{,}085.71$

Next, we focus directly on premium calculation methods. The first and most basic premium calculation is based on the "**equivalence principle,"** which assumes "the net premium is set such that the expected present value of the future loss is zero at the start of the contract" (Dickson, 2013, p. 147). This assumption is mathematically denoted as: $E[L_0^n] = 0$. This directly implies that the expected present value of benefit outgo is equal to the expected present value of collected premiums. This relationship can be used to solve for the **net premium.** Let's revisit our previous example: suppose an insured purchases a one-year $15,000 death benefit life insurance policy for a single premium, P. Assuming the payment will come at the end of the year if the insured dies, and the insured has a 2% risk of dying within the year, we can now calculate the **net single premium** for the policy:

$$EPV(Benefits) = EPV(Premiums)$$

$$(2\%) \times \frac{15{,}000}{1.05} = P$$

$$285.71 = P$$

Now, we will discuss the second key premium calculation method, which uses the **portfolio percentile premium principle (PPP),** sometimes called "Value-At-Risk" (Dickson, 2013, p. 163). The key idea with this second pricing method is that it reflects not only the expectation of the net future loss random variable, but also the variance associated with this risk. The basic question answered by the equivalence principle is: "if I sold a large amount of policies like this, how much would I have to charge each customer to neither make nor lose money, on average?" However, the basic question answered by PPP is: "if I sold a large number of policies just like this one, how much would I have to charge each customer in order to have, at maximum, a 5% chance of losing money on the portfolio?" To answer the latter question, we first assume that we are pricing not one policy, but rather a large portfolio of identical and independent policies. Here, the term 'identical' means that all policyholders will pose the same risk profile. Every policy sold will have its own future loss random variable, $L_{0,i}$, with all N policies having the same mean and variance. 'Independence' means that each $L_{0,i}$ has a correlation coefficient of 0 with

all other policies – so that the outcome of any one policy will not provide new information on the outcome of any other policy. From these assumptions, we can model the behavior of the loss random variable associated with the entire portfolio thusly (Klugman, 2012):

$$L = \sum_{i=1}^{N} L_{0,i}$$

$$E[L] = \sum_{i=1}^{N} E\left[L_{0,i}\right] = N \times E\left[L_{0,i}\right]$$

$$Var[L] = \sum_{i=1}^{N} Var\left[L_{0,i}\right] = N \times Var[L_{0,i}]$$

While a loss chance of 5% was assumed in the example question posed in previous paragraph, this chance is in fact arbitrary. One could just as often wish to price insurance with, for example, an 8% chance of having a positive loss on a portfolio. For the PPP method, we denote this risk measure using the Greek letter 'alpha,' where $1 - \alpha$ is the chance that a portfolio loss will be positive, signifying a financial loss for the insurer ($\alpha = 0.95, \alpha = 0.92\ respectively$). With this notation, we may now price insurance under our new assumptions by beginning with the statement "the portfolio loss random variable should have an alpha-percent chance of being negative" (Dickson, 2013):

$$P[L < 0] = \alpha$$

For large N, we apply the central limit theorem, implying that:

$$L \sim Normal(\mu = N \times \mu_{policy}, \sigma^2 = N \times \sigma_{policy}^2)$$

where μ and σ are now constants pertaining to the losses on the entire portfolio. We then standardize the distribution to find that:

$$P\left[\frac{L-\mu}{\sigma} < -\frac{\mu}{\sigma}\right] = \Phi\left(-\frac{\mu}{\sigma}\right) = \alpha$$

$$\frac{\mu}{\sigma} = -\Phi^{-1}(\alpha)$$

Where Φ is the cumulative normal distribution function. From here, we must write μ and σ in terms of P to solve for P. The exact relationships drawn here will depend on the exact question being posed. Let's borrow from our previous example, where we supposed each insured purchases a one-year $15,000 death benefit life insurance policy payable at the end of the year of death, valued at a 5% interest rate, for a single premium, P. If each insured has a 2% risk of dying within the year, independently from each other insured, and we wish to sell 100 policies such that we have a 5% chance of losing money, then (Dickson, 2013):

$$\mu = 15{,}000 \times (2\%) \times \frac{100}{1.05} - 100 \times P$$

$$\sigma = \sqrt{100 \times \left(\left(\frac{15{,}000}{1.05} - P\right)^2 * .02 + (-P)^2 \times 0.98\right)}$$

The last few steps to solve for P will require standard algebraic manipulation, but the precise solution is nonetheless beyond the scope of this paper.

In summary, LTAM's primary two premium calculation methods are performed under the Equivalence Principle (EP) and the Portfolio Percentile Premium Principle (PPP). Note how although the EP premium is easily calculated, it does not reflect any risk measure beyond the first moment (mean) of a loss distribution. While premiums calculated under the PPP method do account for the second moment of the loss distribution, the calculation is far more complex. Further, the PPP representation of loss volatility is limited to a singular point of focus, the alpha-percent chance of negative portfolio loss. For example, the premium calculation is blind to the expected gain on the portfolio given that there is a gain, and the expected loss conditional on the loss being positive (financial loss for insurer). In addition to these drawbacks, it also requires the limiting assumptions that all policies are both independent and identically distributed, which as previously discussed, is known to be unrealistic for many insurance applications. Finally, the entire calculation process rests entirely on the risk parameter α. While regulators may essentially enforce a maximum value for α, an insurers' choice of α can be arbitrarily lower than mandated, and even the regulator's cutoff itself may be entirely arbitrary in nature. It is then difficult to

objectively infer an appropriate alpha for a given portfolio, in practice. These glaring weaknesses in the academic models taught in LTAM serve as the basis for the problem this paper aims to solve.

2.2.2 Modern Portfolio Theory

The model this paper introduces will require basic knowledge of modern portfolio theory, also called **mean-variance portfolio analysis.** This method was first created in 1952 by Harry Markowitz, who later was awarded the Nobel Prize for his work. The reason this method is so important is because it quantified and confirmed what investors knew for years to be true – that diversification can help an investment portfolio balance its risk. The first assumption of the method is that investors care about investment volatility, and that they are generally averse to it (Markowitz, 1952). Further, investors want two things from their investment portfolio: a high expected return and a low return variance. Thus, for every investor, we may construct a theoretical curve of indifference, which outlines the set of all combinations of return and variance for which an investor is indifferent in investing between each combination. To the left, we see such a curve drawn on independent axes of expected return and "risk" or variance (Fabozzi, 2012). Note that as the level of risk increases, the investor wants a higher expected return in order to feel like the investment provides her the same utility. The only way to increase the utility a portfolio provides is to move up and/or to the left of the curve, thus either increasing the expected return, decreasing the risk, or both. While the exact formulation of such a curve may be difficult, the concept is useful in demonstrating the payoffs between risk and expected return. The real power of this theory is in

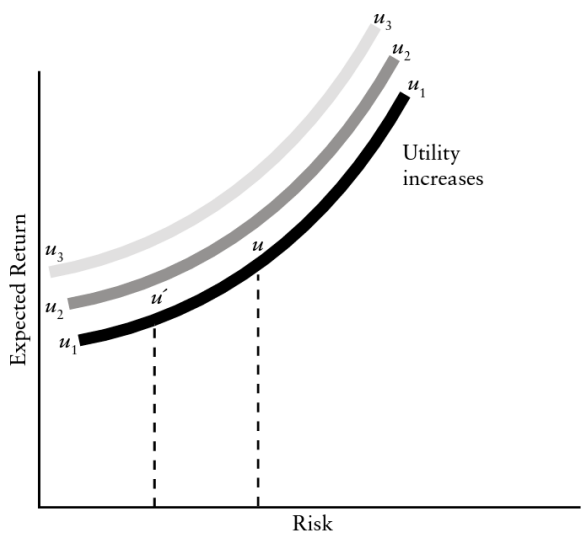

Figure 1: Indifference Curves (Fabozzi, 2012)

shifting the focus of asset valuation from an individual asset to the effect it has on a pre-existing portfolio of assets. The statistical properties of the relationship are trivially calculated from the mean and variance of the returns on the pre-existing portfolio, as well as the mean and variance of the return on the additional asset, and finally the correlation of returns between the asset and the portfolio. This focus on the portfolio and the marginal benefit an asset provides to the portfolio will underpin the conceptual framework of the model laid forth in this paper.

2.2.3 Monte Carlo Methods

The general insurance pricing model set forth in this paper requires basic knowledge of Monte Carlo valuation methods. The SOA's IFM curriculum references the official text *Derivatives Markets* to introduce the concept of Monte Carlo valuation. While it does this from the context of pricing derivatives, it is well suited to pricing various financial instruments. The method generally crops up in finance whenever a closed-form analytical solution (such as those given by EP and PPP for insurance pricing) may not be easily obtained. This is a common occurrence among stochastically defined random variables which exhibit path-dependency. In the context of pricing options, this occurs when the payoff depends on the precise movements of stock prices, hence when "the payoff is path-dependent" (McDonald, 2013, p. 573). The Monte Carlo approach to valuing such methods is to "simulate future stock prices, and then use these simulated prices to compute the discounted expected payoff of the option" by running many simulations and taking the average result across all runs. Assuming the underlying model is unbiased, then by the law of large numbers, we would expect the average of all these simulations to be an unbiased estimator of the stochastic random variable's true expectation. Generally, Monte Carlo valuation is performed under risk-neutral assumptions. This means that the risk-free rate of interest is used in discounting, and that we care only to estimate the first moment of our variable. The key advantage of Monte Carlo valuation is that "with Monte Carlo you simulate the possible future value of the security;

therefore, as a byproduct you generate the distribution of payoffs" (McDonald, 2013, p. 573). In other words, whereas the previous two pricing methods provide limited information about only the moments of the loss random variable, the Monte Carlo valuation can be used to describe the entire distribution of the loss random variable through simulation. This fact heavily inspired the model design featured in this paper.

This concludes the literature review section of this paper. In summary, insurance pricing is a complex process that requires wide breadth of information to price even the most basic products in a realistic setting. Among the key considerations are model precision, loss volatility, enterprise risk tolerance, and assumption flexibility. LTAM's manual sets forth two primary insurance pricing methods -- the Equivalence Principle method and the Portfolio Percentile Premium Principle method. Each method has its own advantages and disadvantages. However, together these methods are not exhaustive in reflecting the aforementioned key considerations of insurance pricing. To meet these business needs with knowledge gained from a standard undergraduate education, a new model will be constructed which is built upon the concepts of mean-variance portfolio theory and Monte Carlo valuation.

Chapter 3

The Model

This section focuses on the original work of the paper, the educational general insurance pricing model. First, detailed assumptions will be made to clarify the context of the model. Then, the model's basic framework will be laid out, accompanied by step-by-step sample calculations to further reinforce the model's underlying process. Next, practical example results will be compiled and analyzed to assess various behaviors of the model. These behaviors will be skeptically evaluated to determine whether they are characteristic of insurance pricing in a real-world setting. Finally, the model's practicality will be discussed with a description of the model's foremost strengths and weaknesses.

3.1 Model Assumptions

The pricing model created in this section of the paper is first and foremost a 'marginal' pricing model. This means that it evaluates the costs and benefits associated with adding one additional policy (or set of policies) to an existing policy portfolio, using basic mean-variance portfolio theory assumptions. Another key assumption the model makes is about the financial objective of the company. As previously discussed, a company might have various competing objectives when pricing insurance. This model assumes that an insurance company has one sole financial objective measurement, the **expected future equity** of the company, previously written as the "Actuarial Present Value" (APV) of future cashflows in the abstract of this paper. Because the sole objective is the expectation of future equity, there is no concern for the variance of future equity. This implies a long-term risk-neutral approach to the insurance pricing model. Since the longest-term investment commonly sold today has a term of 30 years (i.e. a 30-

year US Treasury Bond), the model will assume 30 years is an appropriate planning horizon for all long-term investment planning. As such, the sole measure of financial performance in the model will default specifically to the expected future equity of the company exactly 30 years from present day. The **upfront marginal premium** is the additional premium paid today at which the expected future equity of the company is the same with the new policy as without, given that the policy may already have an arbitrary annual premium associated with it. Moreover, the model solves for this upfront premium by taking the difference between the expected net present value of the company portfolio before and after adding in the new policy.

Thus, to calculate the premium associated with any new policy, the model must first calculate the expected net present value of the pre-existing company portfolio, and then recalculate it with the addition of the new policy. Because the model views company equity as a stochastic process (since tomorrow's equity is a random movement from today's equity), and because the model enforces a permanent state of bankruptcy when company equity bottoms out at zero, the future equity of an insurance company is seen as a path-dependent stochastic random variable. Thus, a Monte Carlo valuation is used to simulate its distribution. The equity movements of a company are modeled annually under the assumption that the portfolio has a set of policies which pays a total fixed annual premium and has a randomly distributed annual loss, which measures the costs associated with annual benefit outgo. Additionally, a risk-free rate of return is earned annually on the previous year's equity holdings, and company equity will stay at 0 forever if it ever crosses zero. The simulated future equity is then described by the following recursive definition:

$$y_t = \begin{cases} MAX[y_{t-1} \times (1+r) + R - l_t, 0] & for\ y_{t-1} \neq 0 \\ 0 \quad for\ y_{t-1} = 0 \end{cases}$$

Where y_t represents the company equity at time t, r is the annual risk-free rate of return earned on equity, R (for revenue) is the annual premiums collected from the policy portfolio, and l_t is the accumulated annual loss amount over the tth year. Specifically:

$$l_t \sim Truncated\ Normal\ (\mu, \sigma^2, truncated\ below\ at\ 0)$$

This annual loss distribution was chosen for two reasons. First, due to the central limit theorem, a suitably large portfolio would expect to have normally distributed annual losses, under the same assumptions that underlie the PPP principle. Second, policy benefit outgo should never be negative, since policyholders filing claims will never pay the insurer to do so, hence the truncation at zero. For our stochastic simulation, we will assume that y_0, r, R, μ, and σ are given constants.

The concept of the model hinges on its ability first to simulate the portfolio before adding the additional policy, and then to simulate it again once the policy is added into the portfolio. All that needs to be done for this is to adjust the loss distribution parameters accordingly. The adjustment is given below:

If $l_{policy} \sim N(\mu_{policy}, \sigma^2_{policy})$, $l_{portfolio} \sim N(\mu_{portfolio}, \sigma^2_{portfolio})$, and $l_{portfolio} + l_{policy} = l_t$, then: $l_t \sim N(\mu_{policy} + \mu_{portfolio}, \sigma^2_{policy} + \sigma^2_{portfolio} + 2 \times \rho_{policy,portfolio} \times \sigma_{policy} \times \sigma_{portfolio})$.

To reinforce the previous section, a very brief example will be reviewed below. If today's company equity is valued at \$50, the risk-free rate of return is 3%, and the annual premium revenue is \$20, let's calculate next year's equity if (a) the company experiences an annual loss of \$30, and (b) the company experiences an annual loss of \$80.

(a) $y_1 = MAX[50 \times (1.03) + 20 - 30{,}0] = MAX[\,41.50, 0] = \$41.50$

(b) $y_1 = 50 \times (1.03) + 20 - 80 = MAX[-8.50, 0] = \$0.00$

Note that in scenario (b), all future equity past y_1 will also equal 0, to reflect a permanent state of bankruptcy.

3.2 Mathematical Analysis

Although the pricing model will be practically solved using a Monte Carlo simulation, it is worth noting that an analytical solution exists, and in fact the model can be practically solved using a program run on symbolic computation. Since the knowledge base required to understand such an application strays far from the core actuarial skills acquired from an undergraduate education, the application will not be a part of this thesis. However, the basic definition behind the approach is worth studying, as it resembles a discrete-time, uncountably infinite state-space Markov chain. With the previously made assumptions, we can define the entire probability distribution of future equity at any future point in time recursively, by the following:

$$\Pr(Y_1 = y_1) = \begin{cases} \Pr(l_1 \geq y_0 \times (1+r) + R), & y_1 = 0 \\ \Pr(l_1 = y_0 \times (1+r) + R - y_1), & y_1 \neq 0 \end{cases}$$

$$\Pr(Y_t = y_t) = \begin{cases} \Pr(Y_{t-1} = 0) + \int_0^\infty \Pr(Y_{t-1} = y_{t-1}) \times \Pr(Y_t = 0 | Y_{t-1})\, dy_{t-1}, y_t = 0 \\ \int_0^\infty \Pr(Y_{t-1} = y_{t-1}) \times \Pr(Y_t | y_{t-1})\, dy_{t-1}, y_t \neq 0 \end{cases}$$

for $t > 1$.

From here, the remaining probability statements are given either by recursive solution or by the probability density function or cumulative distribution function of the truncated normal distribution. This definition, if solved computationally with a “for loop,” for example, would provide all necessary information for the expected future equity of the company simulated, and thus provide analytically precise prices for additional policies. This has major advantages over a Monte Carlo approach, from which we would expect a price estimate variance subject to computational convergence proportional to the root of the number of simulations, possibly an estimate bias, and perhaps most importantly, a far greater demand for intensive computation.

Chapter 4 Application and Results

Now that the mathematical framework of the pricing model has been established, the paper will review the interactive web application built from R-shiny to showcase the model, and then discuss the primary findings of the application and model.

4.1 Generalized Pricing Tool

To demonstrate the educational value of the model, I constructed a web application in R-Shiny for students to interact with. Pictured below is the homepage of the site (https://michael-gregory-callahan.shinyapps.io/Thesis_Demo/).

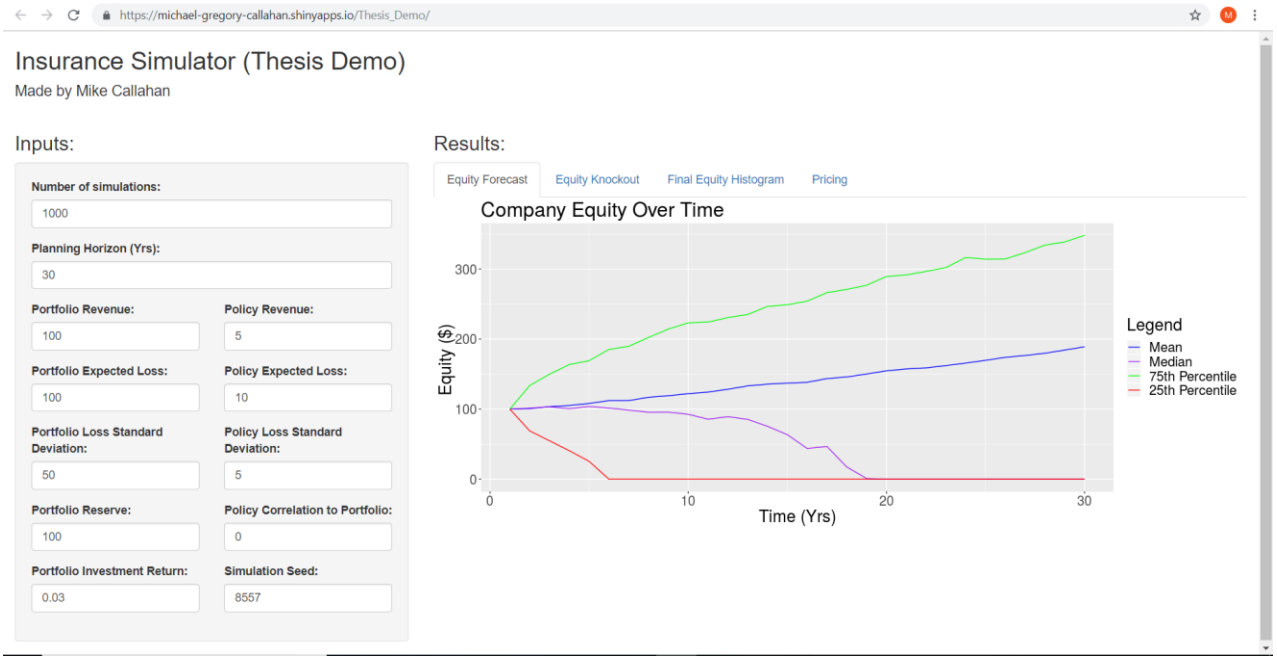

Figure 2: Web Application Home Page

Above, we can see the general layout of the website. To the left, there is a sidebar panel populated with all of the quantitative inputs required to run the model. There, the default inputs are shown for each variable. The user may select any of these white input boxes and change the number inside, and the entire webpage will react accordingly. To the right, we see the main panel, which is programmed as a tabset panel. This means that the user can switch between tabs freely to view different information. It is important to note however, that the first three tabs are displaying information regarding the beginning portfolio, and that the change in risk profile associated with the additional policy is not reflected in these tabs.

The first tab is called "Equity Forecast." Shown there are the 25th, 50th, and 75th percentiles and the mean of future equity over every period simulated. In the above picture, you can see that roughly 25% of all simulations went bankrupt by year seven, and 50% went bankrupt around year 18. However, the mean and 75th percentile lines trend upward strongly.

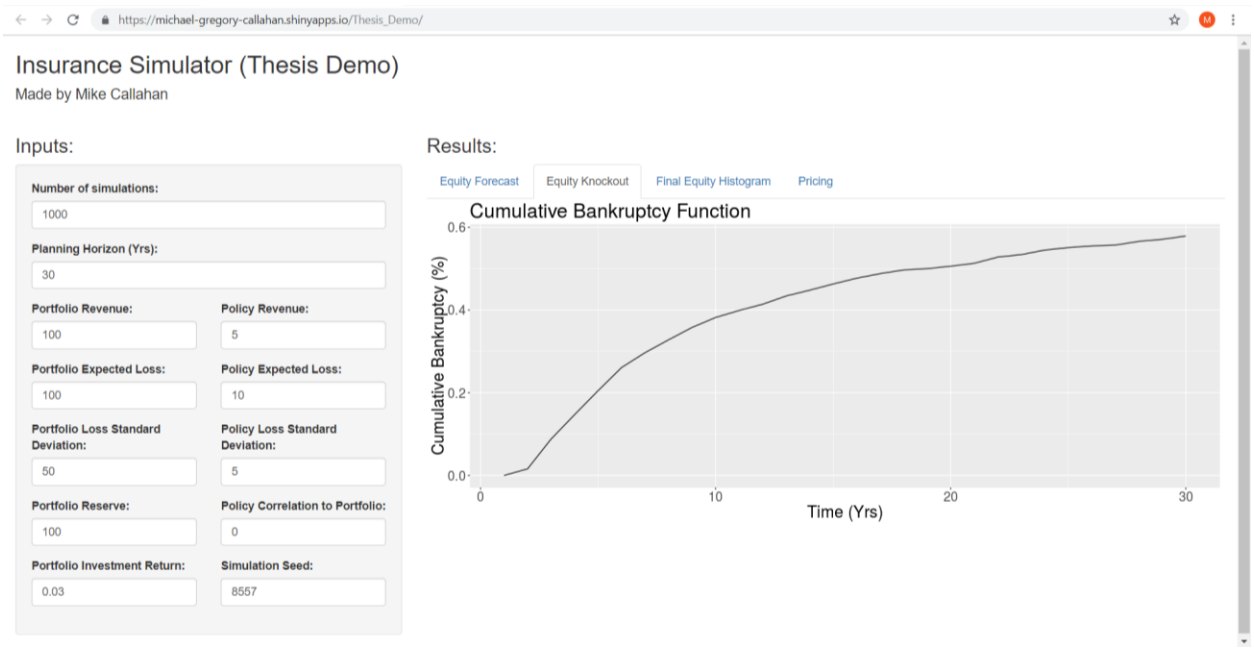

Figure 3: Web App Equity Knockout Tab

The next tab is called "Equity Knockout." This tab, pictured above, provides a graph which displays the cumulative proportion of simulations which have reached bankruptcy over time.

In this particular example, the graph shows that almost 40% of the 1,000 simulations run went bankrupt by the 10th annual iteration, and about 55% were bankrupt by the 20th iteration. Note how the knockout curve bends downward, implying that as time goes on, the likelihood of additional simulations going bankrupt decreases. This makes sense, since the average equity value of simulations which are not bankrupt but the 20th iteration, for example, is higher than at the outset of the simulation. This is an excellent educational example of the importance of conditional expectations. Since interest is earned proportional to the equity value, this means that the portfolio's annual earnings are also higher than at the outset, creating a positive feedback loop.

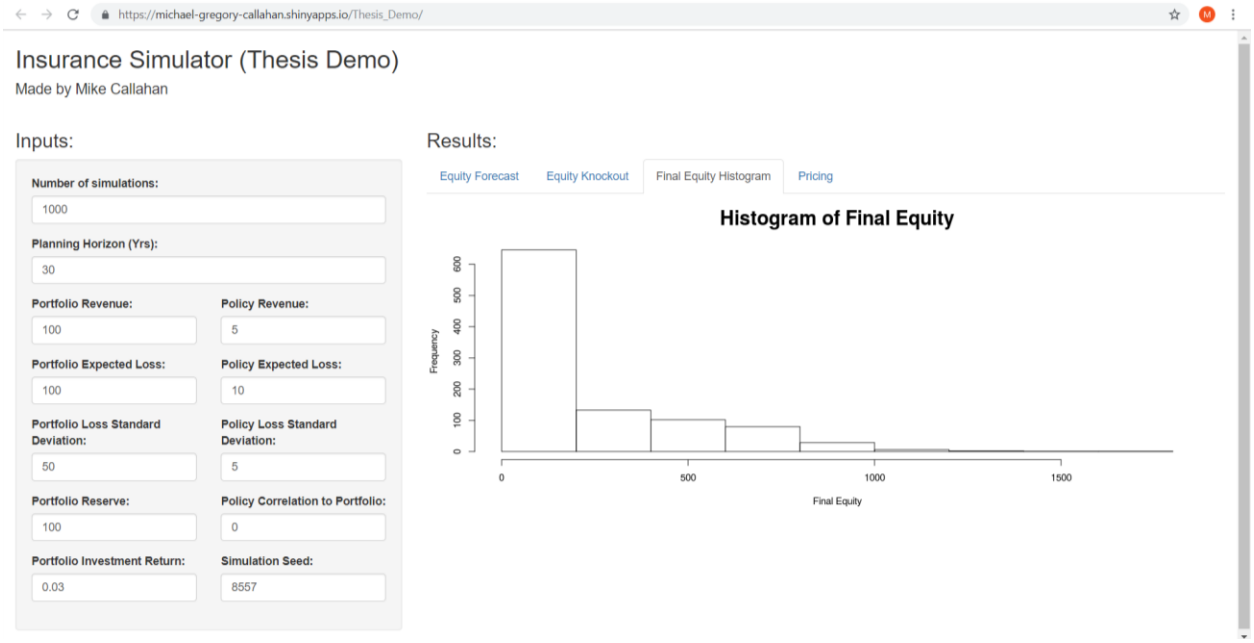

Figure 4: Web App Final Equity Histogram Tab

The third tab is titled "Final Equity Histogram." This tab, pictured above, offers a histogram to display the ending disitribution of company equity across all simulations, which in this case is 30 years from the start of the simulations. Here, we see that over 600 of the 1,000 simulations had less than $200 in equity (most of which were bankrupt) by this time, which is consistent with the ending point of the previously shown knockout curve. Of the remaining simulation, we can see that the final equity value is spread in decreasing frequency from $200 up to $1,000 with a few outliers above $1,000. This confirms

the previous statement regarding the remaining simulations over time, as they have a signifcicantly higher final equity value than the $100 starting reserve. This histogram serves as a rough approximation of the probability density function for the future loss random variable at 30 years.

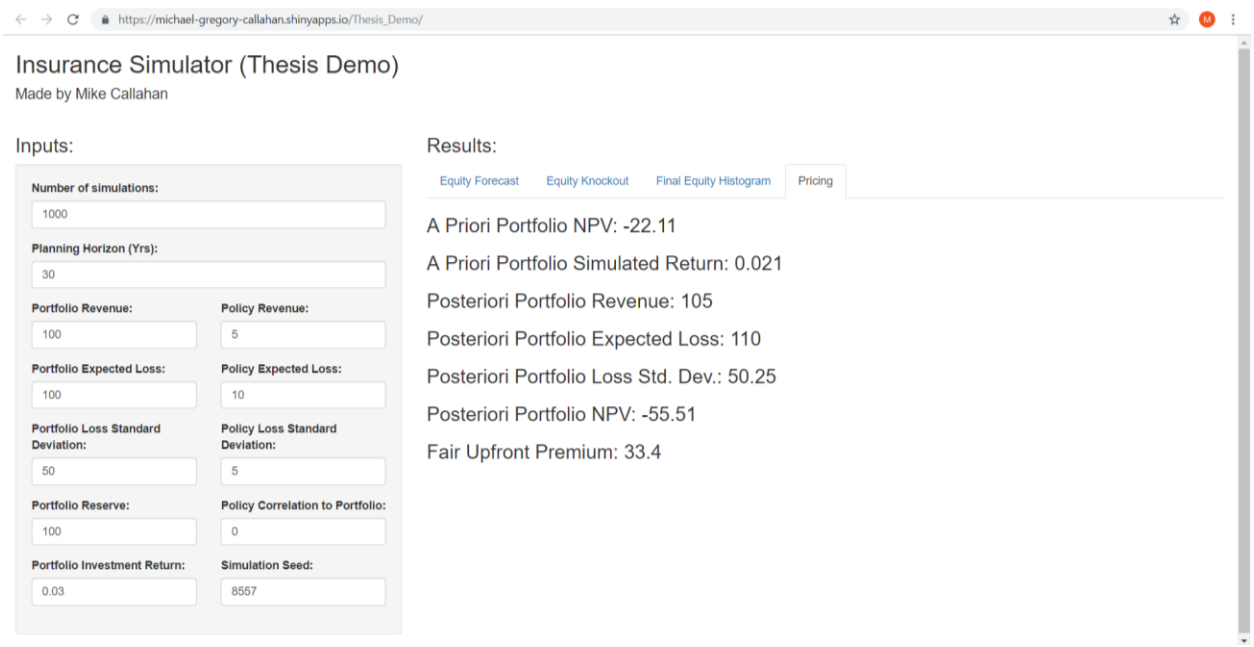

Figure 5: Web App Pricing Tab

The final tab is titled “Pricing.” This tab, pictured above, provides a simple text display of all of the quantitative inputs for calculating the price of the additonal policy. Under “a priori portfolio NPV,” the tab shows that the average final equity of the original portfolio, discounted back to today, less the $100 “initital investment” reserve, is $ -22.11. Thus, the NPV is really showing how the portfolio's performance on that $100 compares to the market rate (set to 0.03 in the simulation). The end number is similar to the discount or premium you might find on bonds.

For example, the "market rate" of 0.03 implies that if this $100 was invested in the market, it would accumulate returns over time, and total $100 \times 1.03^{30} = \$242.73$ in 30 years. But the money is currently invested in the insurance portoflio, which has earnings simulated by the model. Here, the model projects an average final equity of $189.06. This imples that the insurance business provides an

annualized return of 0.021457 (shown above as the rounded figure 0.021), which is lower than the market rate. So the current NPV can be shown as:

$$\frac{189.06}{1.03^{30}} - 100 = \frac{(100\times1.021457^{30} - 100\times1.03^{30})}{1.03^{30}} = -22.11.$$

This means that although you have $100 today for the insurance portfolio, it will earn as much as an amount $22.11 less than if you had invested that $100 in the market. This calculatation is very similar to a zero coupon bond with a face value of $100 being discounted by $22.11 due to rising market rates.

Next, the tab displays how the parameters of the annual loss distribution of the portolfio change once adjusted by the parameters of the loss disitribution of the new policy. Intuitively, we can think of the policy as adding $5 in annual revenue, but costing an average of $10 per year due to losses. So adding the policy is similar to adding in a -$5 cash flow for the next 30 years, on the surface. However, because the equity bottoms out at $0, most of those potential negative cash flows are never realized, which is why the cost will turn out to be a lot less than a $5 30-year annuity immediate. We can see that the NPV of the adjusted portfolio is $ -55.51, which is even lower than the original NPV. This means that adding the policy to the portfolio will lower the portfolio's current market value by $33.40. Thus, the fair upfront marginal price for adding the policy to the portfolio is $33.40.

This concludes the basic review of the functionality of the web application tool I built out in R-Shiny. From here, we will investigate the basic insights such a model has on the pricing of insurance policies.

4.2 Results

Given that the primary objective of the model is to price insurance, it makes sense to study the relationship that the model shows between price and other input variables. Thus, this section will take a brief inventory of the univariate effects of the quantitative model inputs on the output model price. While further investigation may be done to uncover bivariate relationships with price, the number of model variable pairings required, $\binom{10}{2} = 45$, would far exceed the length and scope of this paper.

The default inputs previously shown were not chosen arbitrarily. The projection planning horizon of 30 years reflects the longest-term security investment commonly available today, and the balance of revenue to expected loss and loss variance, as well as the initial reserve, were all chosen to show as wide a range of final equity outcomes as possible while using simple, round numbers. For these reasons, the default inputs will serve as the center of each univariate analysis, and each variable will be adjusted up and down symmetrically from this starting point to determine the variable's isolated effect on price. Further, to limit the computational demand of data acquisition, each variable will only be altered up and down from the base level by 10%, 20%, 30%, 40% and 50%, for a total of 11 data points per variable graph. This should suffice to show the deterministic curves of variable interaction in the model. However, the number of simulations will be increased from 1,000 to 10,000 to allow for cleaner data. One last method of data control is that all simulations run for data acquisition will be run on the same seed, 9,999, meaning that any random model bias will be constant throughout the experiment to further control for random variation in the data.

After this brief inventory is taken, the secondary purpose of the model, which focuses on the effects of policy loss variance and policy-to-portfolio loss correlation on policy price – also known as "risk appetite," will be revisited. The initial hypothesis of this paper is that the risk appetite will vary between risk-seeking and risk-averse behavior, depending upon the financial circumstance of the original insurance portfolio. For background context, an increase in the standard deviation of the additional policy will increase the volatility of the annual losses. Although the expected path of future equity is much the

same, there is effectively a greater chance that the company will go bankrupt, along with a greater chance that the company will have very low losses in any given year.

The prediction is that this balance is asymmetric, and that the asymmetry will vary greatly across different scenarios, even yielding opposing relationships depending upon the exact settings of the simulations run. The intuition for this is as follows: having an exceptionally good year provides lasting advantages through interest on equity; and having an exceptionally bad year would put a stop to all cash flows indefinitely. Thus, when the portfolio is programmed to have a low expected loss ratio, we would expect to it have positive profits in most years. In this case, a per unit increase in volatility should become costlier, because now the bankrupted simulations would have likely been able to come back from bad years, but those positive returns are cut short by bankruptcy. On the other hand, when the loss ratio is much greater than one, the volatility can improve the long-term financial outlook by creating more high-profit outliers each year, which provide lasting advantages, and no negative outliers thanks to the bankruptcy cutoff. We will return to take an inventory of this pattern towards the end of this chapter.

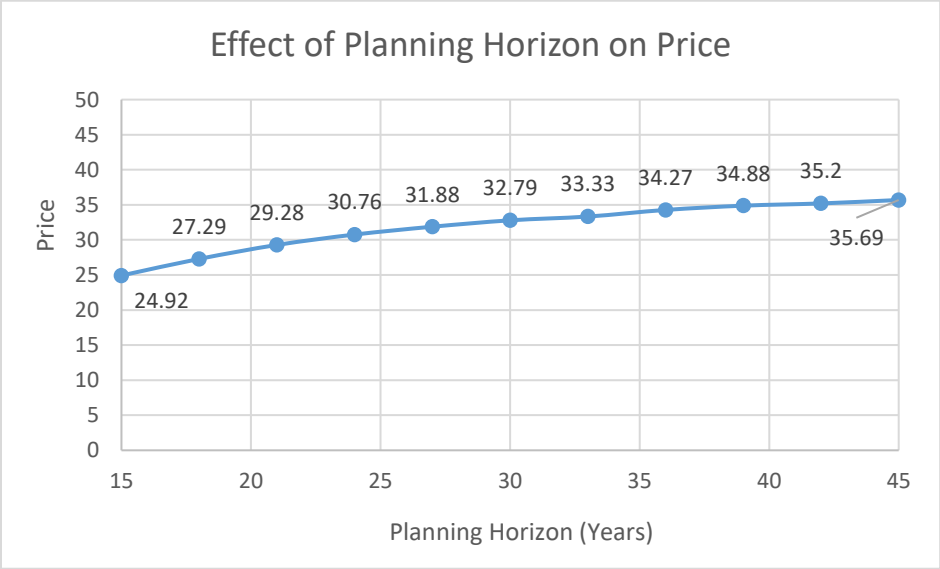

Figure 6: Univariate Effect of Planning Horizon on Price Graph

The first input variable studied was the planning horizon. This variable represents the future point in time and which equity is projected and made the pricing objective of the model. The graph (pictured

above) reflects a very mild positive effect of the planning horizon on policy price. Additionally, the curve seems almost asymptotic, meaning that as the planning horizon increases to a sufficiently large number, the price converges as well. This makes intuitive sense, since changing the planning horizon form 30 years to 100 years will likely have little effect on the cumulative total knockouts, as supported by Figure 3. Additionally, the direction of the curve is intuitively explained by the same phenomenon: less time for the projection results in fewer relative knockouts due to adding in the new policy, hence the policy has less downward effect on the portfolio NPV for shorter time frames and thus costs less.

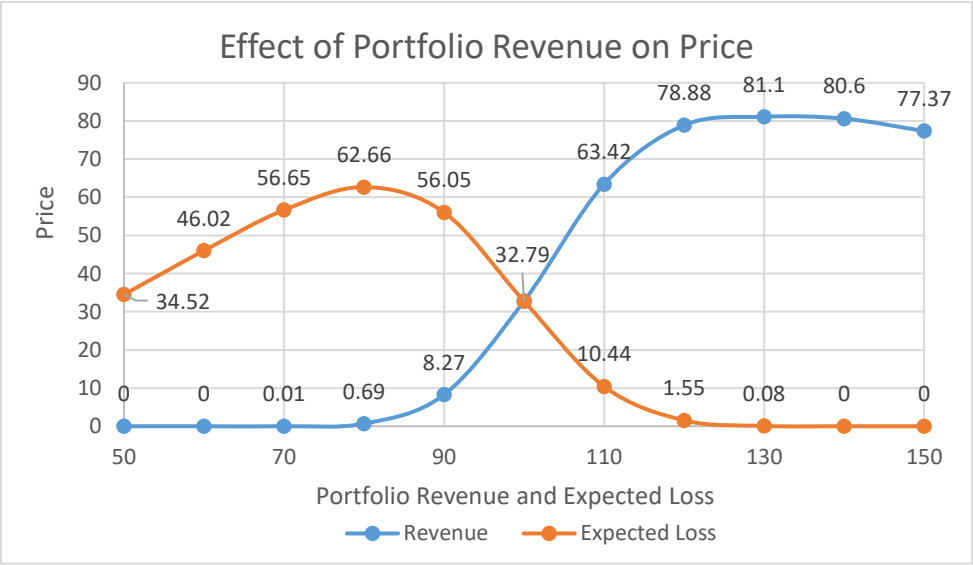

Figure 7: Univariate Effect of Portfolio Revenue on Price Graph

The next two variables studied were the revenue and expected annual loss on the original portfolio of the insurer. These variables yield perhaps the most interesting results, due to their parity and magnitude of effect, and the fact that the previously reviewed academic pricing models do not reflect them in pricing policies. The graph pictured above (Figure 7) demonstrates just how influential the pre-existing insurer loss profile is under the new pricing scheme. As portfolio revenue falls from the base assumptions, or as expected loss rises from the base assumptions, the price of the additional policy falls asymptotically to zero. This is a result of the insurer's certain and immediate financial demise as the result of exceedingly high annual loss ratios. Simply put, the insurer will go bankrupt regardless of the policy

being added in, so the marginal price is effectively negligible. On the other hand, when the revenue is increased, or the expected loss decreased, from their respective base levels, the price of the policy increases to a point, and then comes back down. This pattern is far more complex. Initially, the decreased loss ratio implies less risk of future bankruptcy in the original portfolio, and thus greater relative increase in bankruptcy risk once the new policy is added, which bends the NPV down more and increases the price of the new policy. However, after a certain point, perhaps the policy's effect on the likelihood of bankruptcy wears off as the portfolio becomes so profitable that even the new policy cannot greatly affect bankruptcy outcomes over the projection period.

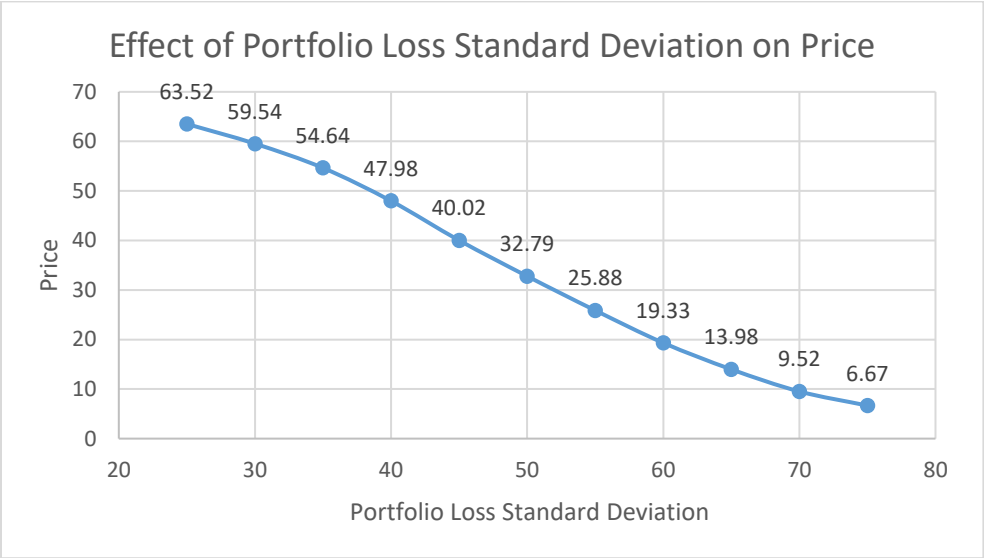

Figure 8: Univariate Effect of Portfolio Loss Standard Deviation on Price Graph

The next variable studied was the original portfolio's annual loss standard deviation. This is another variable which traditional academic pricing models do not account for, yet it is shown in the graph to above (Figure 8) to have a strong impact on the price of the additional policy. The curve has a decreasing logistic shape, with an inflection point at 50. The direction of the curve implies that the higher the initial volatility in the original portfolio, the less the additional policy will negatively affect the portfolio's performance. Additionally, the curve seems to have two horizontal asymptotes, which may serve to be equilibria points across the portfolio loss standard deviation.

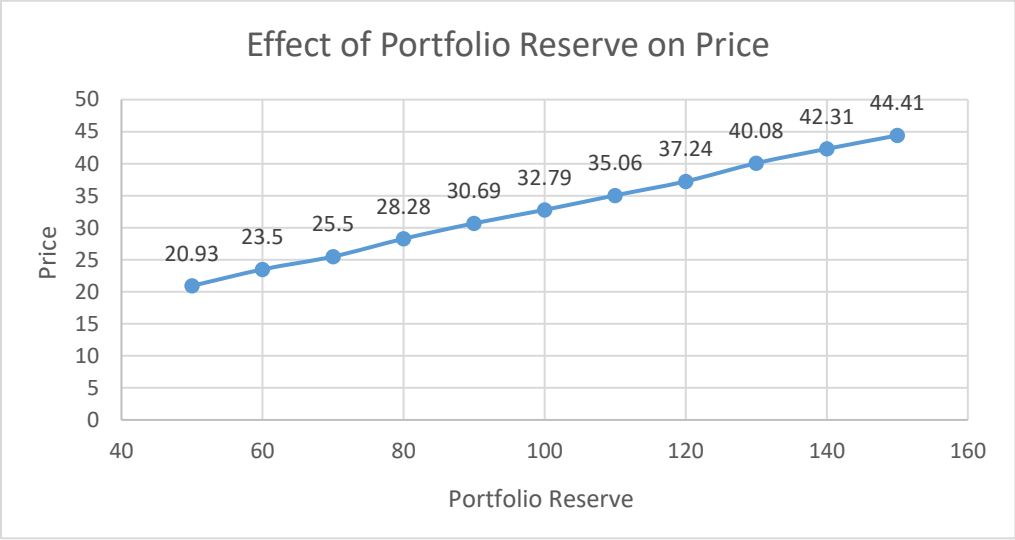

Figure 9: Univariate Effect of Portfolio Reserve on Price Graph

The last variable concerning the original portfolio is the starting portfolio reserve. This is the beginning equity at time zero in the simulations. As we can see from the graph above (Figure 9), the reserve seems to have a direct linear relationship with the price of the policy. This pattern reflects all previously seen patterns in the model: any change in the original portfolio which would increase business longevity -- in this case, increasing the starting equity value -- will make the additional policy costlier to the portfolio's overall performance. Applying the information gleaned from the previous graphs, however, it is unlikely that this curve is linear in nature throughout the domain of portfolio reserve values, and therefore likely are seeing just a 'zoomed-in' snapshot of another, much larger logistic curve.

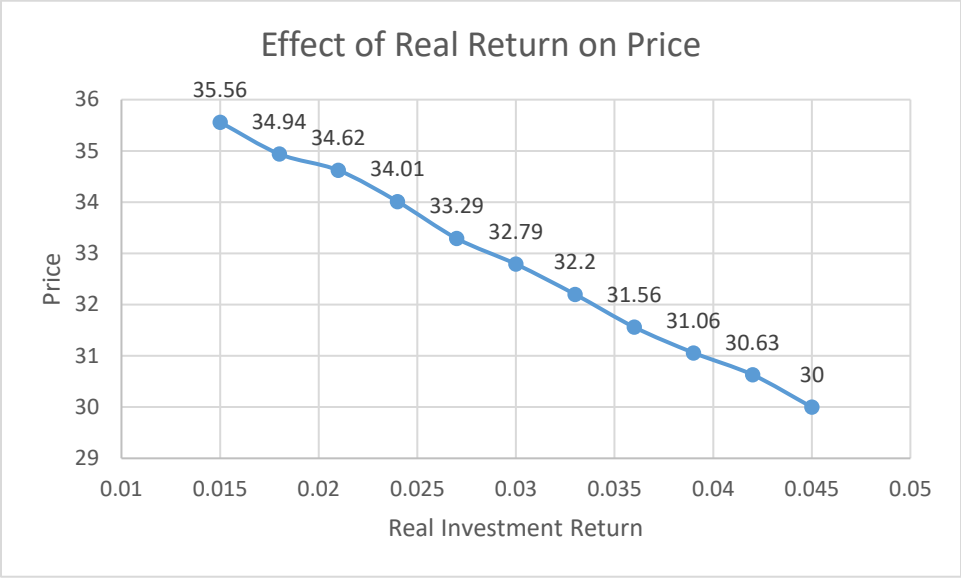

Figure 10: Univariate Effect of Real Return on Price Graph

The next variable studied was the real return on annual investments. This variable determines the annual income earned on the previous year's equity holdings. The graph above clearly displays an approximately linear curve which shows a mild inverse relationship with policy price. This is graph is fascinating because it breaks the previously observed pattern, whereby variables moving in the direction of greater original portfolio profitability and sustainability were associated with upward movements in price. In this case, the greater the income earned, the less costly the additional policy becomes to the portfolio's performance. One likely explanation for this is that the additional policy's effect on cash flows has a duration (or cashflow-weighted average time) similar to that of the original portfolio, and thus since the NPV of the portfolio is made more negative by the addition of the policy, the larger loss (that with the policy) is discounted down in magnitude more than that of the original portfolio's loss. For example, suppose the original portfolio had an NPV of - $30.00, and the additional policy decreased this NPV by $20 to have a new NPV of $-50.00. If we subject both NPV's to the same additional discount of 20% (which would be a supported assumption under the equivalence of durations between portfolios), then we

would arrive at a new NPV of -$24.00 for the original portfolio and -$42.00 for the adjusted portfolio. This would imply that the marginal price decreased from $20.00 to $18.00.

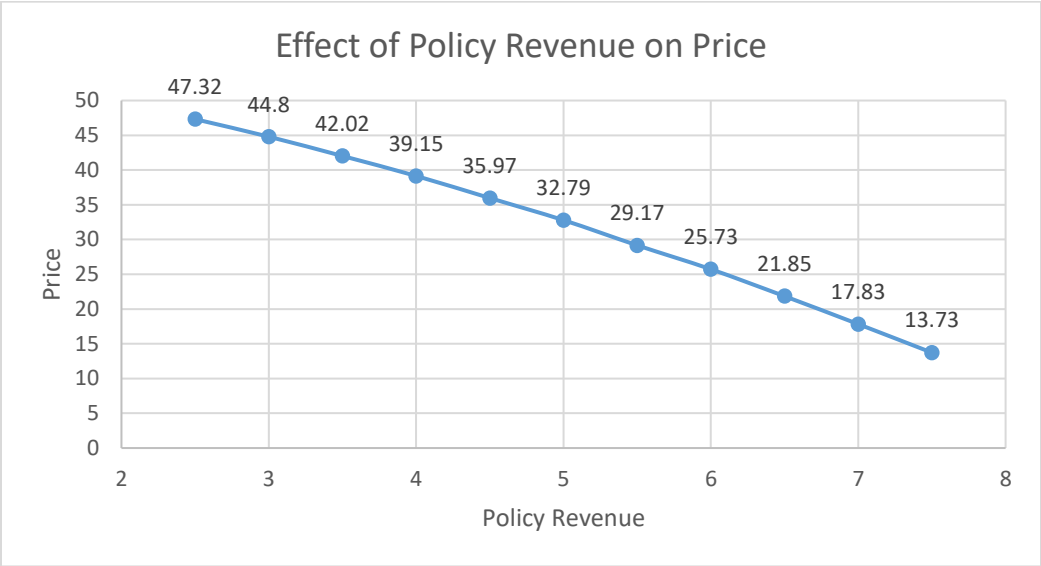

Figure 11: Univariate Effect of Policy Revenue on Price Graph

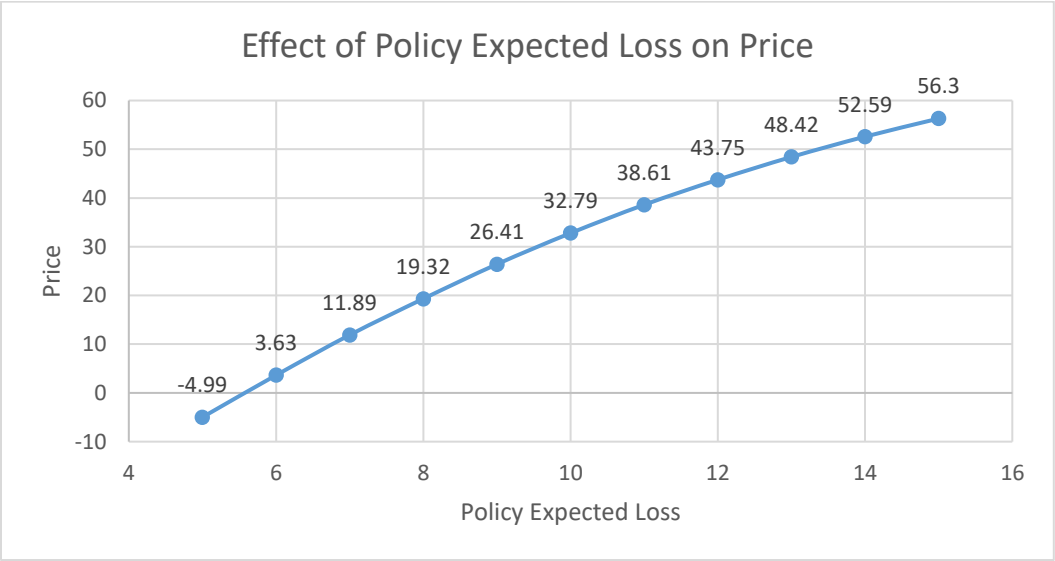

Figure 12: Univariate Effect of Policy Expected Loss on Price Graph

The next two variables pertain to the additional policy being priced. First, the policy revenue represents the annual income coming in from the addition of the policy to the original portfolio, and is ostensibly an annually recurring premium on the policy. The second variable, the expected annual loss on the policy, represents the mean of the truncated normal random variable, which represents annual

losses due to the policy alone. These two variables serve as complements, just as the revenue and expected loss on the original portfolio did. Here in Figures 11 and 12, we see an identical pattern to those two previous variables; however, due to a smaller scale of variation, we see a more 'zoomed-in' snapshot of the same pattern. This is the pattern which would indeed be captured by traditional pricing models, though not in the exact same shape as pictured here, due to a difference in financial objectives.

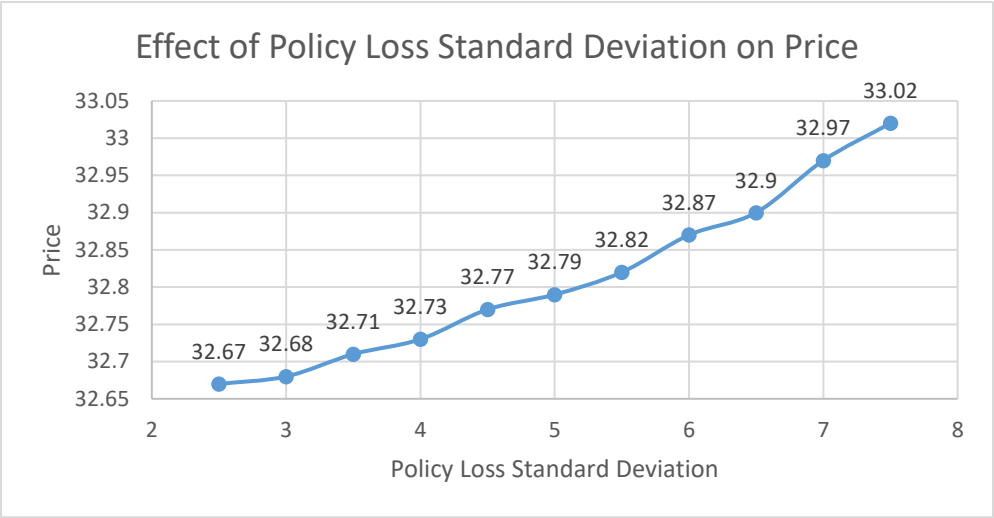

Figure 13: Univariate Effect of Policy Loss Standard Deviation on Price Graph

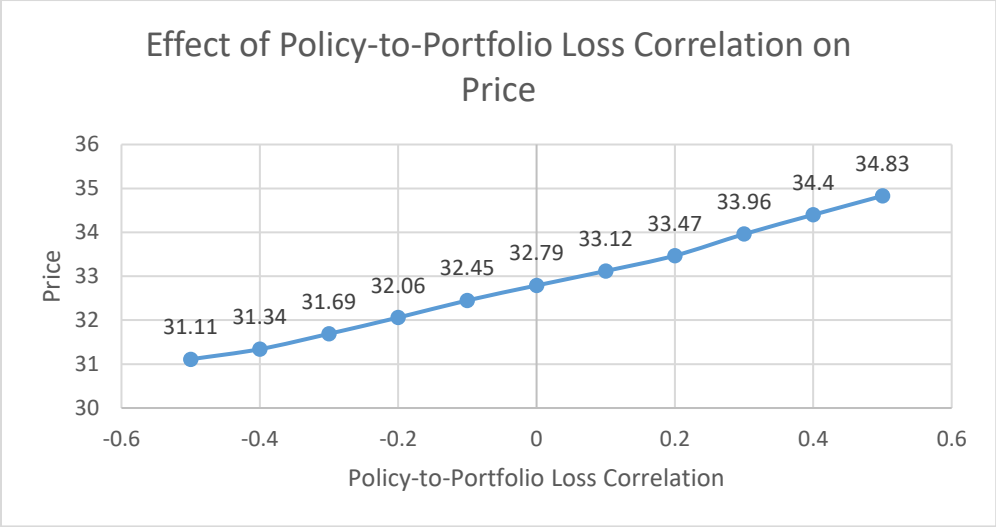

Figure 14: Univariate Effect of Policy-to-Portfolio Loss Correlation on Price Graph

The final two variables of the analysis are the policy loss standard deviation and the correlation coefficient between the policy loss and portfolio loss. These two variables are obviously tied to the standard deviation of the original portfolio in a three-way interaction which determines the overall volatility of the annual loss of the newly adjusted portfolio. As such, these two curves above (Figures 13 and 14) display a similar pattern to each other, because the policy-to-portfolio correlation coefficient is effectively a scaling factor of the policy loss standard deviation. We can observe that the higher the additional volatility of the adjusted portfolio due to the policy, the costlier the policy becomes. This is a direct reflection of how the current portfolio being modeled favors a risk-averse pricing scheme. However, it is worth noting that the model's sensitivity to these variables is very low. This is likely because the base assumptions for this analysis were deliberately chosen to show a wide range of final equity outcomes, so the risk-appetite is in fact more neutral than it would normally be.

	Price			% Change from base price		
Policy Loss Std. Dev.	80% LR	100% LR	120% LR	80% LR	100% LR	120% LR
30	125.01	40.65	1.19	39.38%	13.20%	-13.14%
28	117.13	39.56	1.22	30.59%	10.16%	-10.95%
26	109.46	38.57	1.26	22.04%	7.41%	-8.03%
24	102.39	37.57	1.3	14.16%	4.62%	-5.11%
22	95.97	36.72	1.34	7.00%	2.26%	-2.19%
20	89.69	35.91	1.37	0.00%	0.00%	0.00%
18	84.48	35.29	1.41	-5.81%	-1.73%	2.92%
16	79.47	34.67	1.43	-11.39%	-3.45%	4.38%
14	75.2	34.16	1.46	-16.16%	-4.87%	6.57%
12	71.32	33.63	1.48	-20.48%	-6.35%	8.03%
10	68.04	33.28	1.51	-24.14%	-7.32%	10.22%

Figure 15: Effect of Loss Ratio on Risk Appetite Data Table

To test this idea, the same analysis was performed, but with the base assumption of the standard deviation of annual loss on the additional policy increased to 20 (to magnify the relationships observed), and the expected loss on the original portfolio varied up and down by 20 from the base assumption of 100, to provide a look at risk appetite across different expected loss ratios of 80%, 100% and 120%. Then, the standard deviation was varied up and down 10%, 20%, 30%, 40%, and 50% from the new base

assumption of 20, and the corresponding changes in price were recorded. Finally, the corresponding prices were converted into %-changes from the base price to control for the direct effects of loss ratio on price (Figure 15). Since the correlation factor previously discussed is simply an amplifier of adjusted portfolio volatility, one would expect a price relationship identical to policy loss variance, so the paper will not repeat the experiment for policy-to-portfolio correlation. One glaring detail not discussed in the following analysis is that the prices change dramatically in scale going from a high loss ratio to a low loss ratio. This means that as a company becomes less financially solvent, it becomes less sensitive to additional future losses, and thus more willing to charge less in premium to take on costly policies.

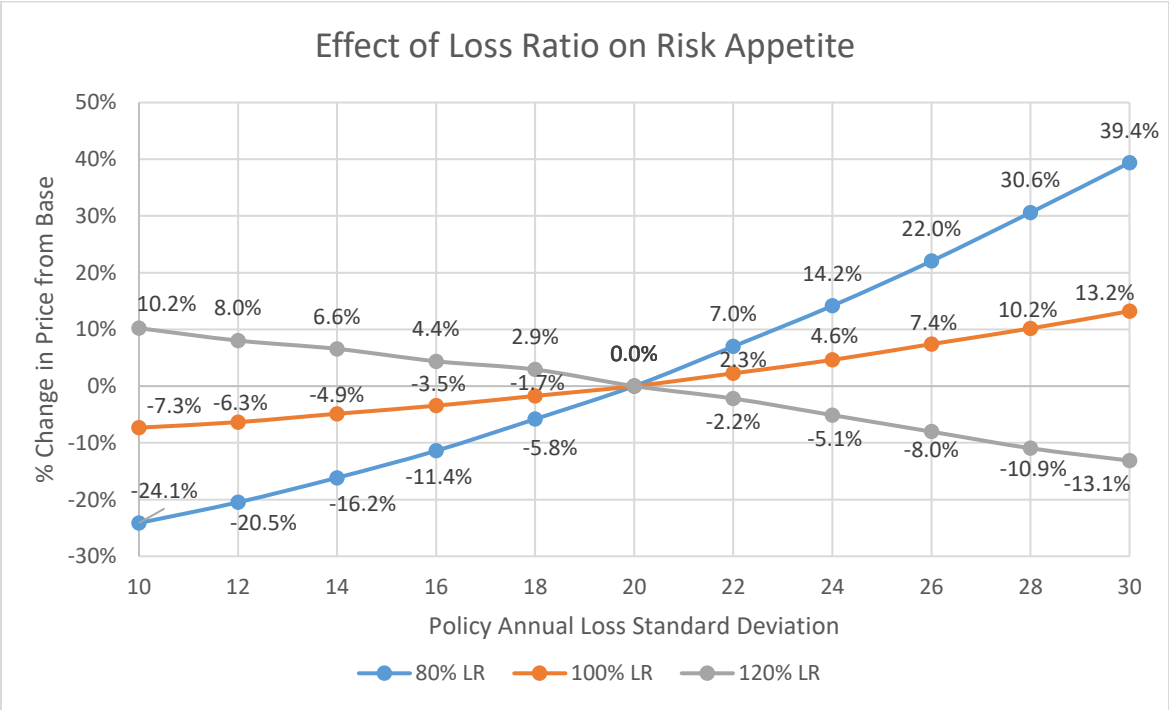

Figure 16: Effect of Loss Ratio on Risk Appetite

The findings of the experiment are presented on the graph above (Figure 16). On the graph, we can clearly see that the relationship between price and policy loss volatility does in fact invert across the different assumptions on the loss ratio of the original portfolio. The 80% loss ratio portfolio put a high price on volatility, resulting in a strong risk premium moving from left to right on the graph; the 100% loss ratio portfolio put a low price on volatility, resulting in a mild risk premium moving from left to right

on the graph; and the 120% loss ratio portfolio put a small bounty on volatility, resulting in a mild risk-based discount moving from left to right on the graph.

Thus, we conclude that the data supports the paper's hypothesis: that in the model, risk appetite is dynamic and can vary across different financial circumstances. From afar, this might seem counterintuitive, as a company identified as a "sinking ship" can thusly be incentivized to engage in even riskier practices to right itself, resulting in the high likelihood of an even faster demise. The paper will therefore label the counterintuitive relationship the "Sinking Ship Paradox."

Chapter 5

Conclusion

This section will evaluate the merits of the proposed pricing model, regarding the conceptual foundations, educational value, and practical advantages and disadvantages, and pricing insights of the marginal pricing model. Reflections will also be made on the effectiveness of the R-Shiny tool, and possible avenues for future investigation.

While the two conceptual models currently used to educate undergraduate students are necessary components to understanding the actuarial industry, they are not exhaustive in providing all the basic insights required for students to understand the most common pricing processes in industry. Most notably, these pricing models rely on the independence of loss risks and consider only arbitrarily defined cutoffs for mitigating portfolio risk. The marginal pricing model provides a new perspective for pricing general insurance, where no explicit assumptions about risk appetite need to be made. With these oversimplifying assumptions removed, one sees that the marginal price provides great improvements in accuracy, particularly under three different circumstances: first, when the pre-existing book of business is highly volatile; second, when the additional policy's annual loss possesses a high variance in relation to the pre-existing portfolio's loss variance; and third, when the additional policy has some observable loss correlation to the original portfolio of business. However, when these circumstances are turned around, and a vastly larger book of business is compared to a much smaller additional policy, both with relatively low loss ratios, then the marginal price will become very similar to the results obtained from simpler pricing methods. This implies that the marginal pricing model

might be a worthwhile effort for more boutique insurance companies, or even reinsurers which possess a small number of very large policies.

However, while the marginal pricing model allows for more parsimonious assumptions and thus a stronger sense of realism over the other two models, the practicality of its application is more burdensome. The model may be solved analytically with a symbolic computation engine, requiring highly advanced software, or computationally, which demands a large of amount of brute-force computational power for acceptable levels of convergence. Both options are difficult to implement, especially for a boutique firm small enough for the underlying theory to still be advantageous. Thus, this is a model which is likely best implemented by reinsurers or developed into a highly advanced analytical software package to be licensed to smaller boutique firms.

This difficulty of implementation reduces the educational value of the model. If students cannot easily apply the model on an exam to assess their knowledge, it will be less likely that the model will be discussed in the first place. However, the model can still be accessed via the interactive web application, which can provide deep insights to students on the nature of insurance pricing and expose them earlier on to a direct application of various actuarial techniques, namely financial forecasting and simulation. Additionally, the students can be made to hypothesize the relationship of the input variables to the output price, and then test those hypotheses in real time. These benefits may be worth the time to further expand on the paper and create instructional materials for educators.

The relationships exposed in this paper's analysis section reveal rich patterns in insurance business strategy. Firstly, the better off a company is, the more money it will demand in upfront premium for a given additional risky policy, and vice-versa. However, this pattern, like most economic patterns, shows diminishing returns in either direction. It also shows how simply

setting one's financial planning horizon further into the future makes one's behavior more risk-averse in the short term. One might argue that this knowledge could be applied to executive compensation to mitigate risky behavior. Further, the analysis suggests that an increase in interest rates effectively shortens this planning horizon by making later future payoffs less valuable, thus increasing the short-term risk tolerance of the insurer. Finally, the analysis shows how crucial the interactivity of loss distributions is when modeling insurance prices under certain circumstances. Risk appetite, for example, can completely invert and then speed up the bankruptcy of a firm via the "Sinking Ship Paradox." These insights demonstrate the need for integrative enterprise risk management in insurance, the dangers of isolating the performance projections of segregated policy portfolios, and most importantly, the need for regulation in the insurance market.

Certain areas of the marginal pricing model warrant further investigation. Firstly, a direct application of the analytical solution would prove to be invaluable in obtaining further insights into the model, especially in relation to long term equilibria and infinite limits. Secondly, within the computational model, with greater resources one could provide wider views of univariate interactions in the model, as well as multivariate interactions, and even measures of computational convergence, to determine how many simulations need to be run for the scalable results to be significant to a single dollar of policy premium. Finally, one might even endeavor to expand upon the theoretical framework that the model provides in several ways: by adding in path-dependent dividend payoffs from annual earnings; or by programming the interest rate as a random variable; or by programming in the market forces of supply and demand.

In conclusion, the marginal prices model provides great advantages over existing models, but only in very specific circumstances. The best use cases for the model are in reinsurance, or

third-party software development for smaller boutique insurance firms. The model can be used educationally, to introduce different actuarial concepts, and the associated tool can be used to test variable relationships in real time. Finally, an analysis of the model's results generated some deep insights into insurance business practices; namely, the "Sinking Ship Paradox."

BIBLIOGRAPHY

1. Cather, D. A. (2018). Cream Skimming: Innovations in Insurance Risk Classification and Adverse Selection. Risk Management and Insurance Review, 21(2), 335-366.
2. Dickson, D. C., Hardy, M., & Waters, H. R. (2013). Actuarial Mathematics for Life Contingent Risks (2nd ed.). *New York: Sheridan*.
3. Fabozzi, F. J. (2012). Encyclopedia of Financial Models. *John Wiley & Sons*. Retrieved from https://ebookcentral.proquest.com.
4. Klugman, S. A., Panjer, H. H., & Willmot, G. E. (2012). *Loss Models: From Data to Decisions*(4th ed.). Hoboken, NJ: *John Wiley & Sons.*
5. Markowitz, H. (1952). Portfolio Selection*. *The Journal of Finance,7*(1), 77-91.
6. McDonald, R. L. (2013). Derivatives Markets (3rd ed.). *Pearson*.

MICHAEL G. CALLAHAN

Academic Vita

Actuarial Exams

Exam P/1: Passed **Exam FM/2: Passed** **Exam MFE/3: Passed** **Exam LTAM/4: Passed**

Education

The Pennsylvania State University, *Schreyer Honors College* **August 2015 – May 2019**
B.S. in *Mathematics* with Actuarial Option, Eberly College of Science, graduated May 2019.
Minors in *Statistics* and *Information Science & Technology*.

- *Presidential Scholar Award; Phi Beta Kappa; Mu Sigma Rho.*
- Economics, Finance and Statistics VEEs completed.

Experience

New York Life, New York, NY **May – August 2018**
Actuarial Intern, Investment Strategy.

- Led "Deep Dive" asset/liability modeling analysis to reinvest $3.2B Long Term Care asset portfolio. Revised existing framework for sensitivity analysis by computing metrics in R to evaluate multidimensional cross-correlated assumption shocks. New framework to be used in all future investment analyses.
- Designed new investment data visualization techniques in Excel and R. Automated graphing process to pull from modeling software. New template will be used to visualize all future "Deep Dive" analyses.
- Created Excel algorithm to automate monthly attribution of $4.5B Fixed Deferred Annuities assets by reconciling a priori and a posteriori investible assets with product sales. Will soon be used across all annuity products.

Chubb Limited, Philadelphia, PA **May - July 2017**
Actuarial Intern, Global Analytics.

- Project Lead for statistical pricing model. Created and presented final model design and dataset to stakeholders two weeks ahead of schedule. Model was approved to begin analysis stage before my departure and will be used to help price $250M of Media Errors & Omissions policy premiums.
- Revised legacy code to update modeling data; performed Defined Book Run for all Q2 financial lines.
- Constructed neural network in Python to model EPL claim severity as a proof of concept. First time department had a working model in Python, and first predictive neural network. Wrote interactive code documentation.

Cold Spring Harbor Laboratory, Cold Spring Harbor, NY **September 2014 - May 2015**
Martienssen Laboratory for Plant Biology; Partners for the Future Program

- Worked with Dr. Rowan Herridge to conduct original, data-intensive research in plant genomics, studying RNA interference and Argonaute proteins to establish gene-backup framework for engineering mutagenic-resistant crops. Won three scholarship awards and Highest Honors Award at science fair with over 3,000 participants.

Involvement

Actuarial Teaching Assistant **August 2017 – December 2018**

- Create course curriculum, homework assignments and exams under faculty supervision for *Risk Management 297C – Introduction to Actuarial Science and Probability Theory*.
- Present weekly lectures to prepare 30+ students for the P Exam; grade assignments, projects and exams.

Undergraduate Research **August 2017 – April 2019**

- Write honors thesis on novel insurance pricing method implementing cross-correlative insured discounting.

Penn State Learning Peer Tutor (Mathematics) **August 2016 – May 2019**

- Tutor 10+ math courses to students individually and in group sessions; create and record online course content.

Technical Skills

Programming: Excel, Mathematica, MATLAB, Prophet (Actuarial Modeling), Python, R, SAS, SQL, and VBA.

"The first principle is that you must not fool yourself, and you are the easiest person to fool."
-- Richard Feynman