MEASURING AND MODELING MARKET RISK FOR LIFE INSURANCE COMPANY ASSETS: AN APPLICATION OF EXTREME VALUE STATISTICS

By

Ryan Timmer

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Business Administration – Finance – Doctor of Philosophy

2015

ABSTRACT

MEASURING AND MODELING MARKET RISK FOR LIFE INSURANCE COMPANY ASSETS: AN APPLICATION OF EXTREME VALUE STATISTICS

By

Ryan Timmer

Standard deviation and variance have been the default measures of investment risk at least since Markowitz's seminal contribution to portfolio selection in 1952. Intuitively, though, investors may not be symmetric around the mean in their attitude toward risk. In other words, they may be much more concerned about the possibility that realized returns are significantly lower than expected ("left tail risk") than the alternative of returns being significantly higher than expected. We study the asset allocation decision for a life insurance company, which is an environment where left tail risk is of utmost concern to the investor. Due to the long-term nature of a life insurance company's liabilities, the insurer must necessarily select asset portfolios with a high premium on avoiding left tail risk for regulatory and long-term profitability reasons. We use extreme value theory, downside risk measures, and copulas to model the market risk of a life insurer's assets for the purpose of selecting an optimal portfolio in such an environment. We find that the current industry allocations to at least one of the primary drivers of life insurer market risk (equities) are close to optimal as of 2013. In addition, we study how the optimal General Account corporate bond and equity allocations, which are chosen by the company, are affected by policyholder investment decisions in the Separate Account and other allocations in the General Account over the past two decades.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

I. Introduction

A. Role of Life Insurance in Society

Life insurance is an extremely important part of our financial system and economy. Like other financial institutions such as commercial and savings banks, pension funds, investment management companies, etc., life insurance companies engage in the economically important activities of the financial system. One can think of finance as being the heart of the economy. As the heart pumps blood out to the different parts of the body making sure each part receives the oxygen it needs and only what it needs, the financial system pumps the economy's lifeblood to where it is needed. When functioning properly, it allocates capital from those willing to supply it (the savers) to those who need it (the borrowers). By doing so, it directly benefits the savers and the borrowers by bringing them together as well as enabling a well-functioning economy.

In playing its part in the financial system, the life insurance industry supports two major components of the economy. The products of this industry, which are mainly life insurance and annuities, provide financial security and stability to 75 million households in the United States as of 2009 (Ernst & Young (2014)). Ernst & Young has also estimated that of these households with life insurance coverage, such coverage would provide sufficient resources to 82% of the children in these households to maintain their current standard of living for a year. In contrast, the financial assets of only 4% of the children in households without life insurance coverage would be sufficient to maintain their current standard of living for a year. Annuities can also provide significant benefits to the personal finances of retirees by ensuring they will not outlive their source of income. Perhaps these are some of the reasons that at least two life insurers have been named systemically important financial institutions by federal regulators.

Life insurance companies pool together significant amounts of capital by collecting premiums from their policyholders that must then be invested at stable returns to support the terms of the policies. As a result, this industry also supports the corporate and government sectors of the economy by providing much-needed capital. In fact, the industry collectively finances about 20% of the corporate and foreign bond market in the United States and about 12.5% of the commercial mortgage market as of the end of 2012 (Board of Governors (2014), Tables L.212 and L.220). Clearly, a financial crisis in the life insurance industry would have widespread and detrimental effects on many other economic actors. Households would have a weakened safety net when they are at some of their most vulnerable times financially, and many businesses would lose a key source of financing.

B. Literature Review and Our Contributions

In light of the life insurance industry's importance as a bulwark against financial storms for policyholders and as a provider of capital to the economy, it seems pertinent to study the asset allocation choice faced by life insurance companies. In addition, there appears to be great importance in studying how this choice ought to be made in light of the dire consequences to both the owners and the broader economy should companies fail to have sufficient resources to fulfill the long-term promises made to policyholders. This is the fundamental question we seek to study here. How should this important part of the financial system, which has not thus far received much attention in the finance literature, approach its asset allocation decision?

Asset allocation as a concept is not a novel contribution of modern finance. We know that the importance of how one allocates wealth to individual investments has been recognized in some form for many centuries. The Babylonian Talmud contains the following 1,500-year-old advice, "A man should always place his money, one-third into land, a third into merchandise and keep a third in hand" (Levy and Duchin (2010)). Even earlier comes this anecdote about Jacob, a patriarch of Israel, "…he divided the people who were with him, as well as his flocks, herds, and camels, into two camps. 'If Esau should attack and overwhelm one camp,' he reasoned, 'the remaining camp may still survive'" (Holy Bible, Genesis 32:8b-9). More recently, Shakespeare writes in *The Merchant of Venice* that Antonio takes comfort in knowing that, "My ventures are not in one bottom [ship] trusted, / nor to one place; nor is my whole estate / upon the fortune of this present year…" (Shakespeare (1598)) These references exhibit a fundamental principle of diversification that it is unwise to place all of one's "eggs" (wealth) in one "basket" (investment).

Still, Markowitz (1952) provided a revolutionary contribution to our understanding of asset allocation by systematically solving for an optimal distribution of wealth across one's portfolio by mathematically optimizing the inherent risk-return tradeoff. Since his seminal papers on this topic, risk measurement for the sake of portfolio construction has largely been based on his choice of variance and standard deviation. However, work by Roy (1952) provides an alternative way of thinking about risk. He approached the asset allocation problem with an approach that risk is the potential for disaster or catastrophe to occur. Thus, it is focused on the most extremely negative deviations from expected returns rather than both positive and negative deviations. Hence, his asset allocation approach has been referred to as "safety-first" in the sense that an investor utilizing his approach is seeking to maximize return while minimizing the chance that ruinous outcomes occur. This risk is also referred to as tail risk because the left tail of the probability distribution of returns is where these outcomes reside. Although Roy's work has received much less attention over the subsequent decades, his is potentially the more relevant in the context of an asset allocation decision for a life insurance company. Thus, we will study the asset allocation problem of a life insurance company in a downside risk framework with Roy's "safety-first" view of risk.

Roy's work has been extended to focus exclusively on the asset allocation decision faced by a life insurance company. Browne (1995), Liu and Yang (2004), Chiu and Li (2009), Consiglio, Pecorella, and Zenios (2009) propose optimal investment strategies for an investor seeking to minimize their probability of ruin. This is also in the line of work that has been done to develop asset-liability management ("ALM") models for the case of a life insurance company (e.g., Lamm-Tennant (1989), Sharpe and Tint (1990), Sherris (1992), Consiglio, Cocco, and Zenios (2008), and Chiu and Li (2009)). Although some of this work studies the ALM problem within a downside risk framework, much of it is theoretical in nature. The general dearth of data on life insurance company liabilities makes empirical analyses of these ALM models difficult.

Another limitation of these models is that they often include only one or two asset classes while actual life insurance companies invest in a wider range of asset classes including stocks, corporate bonds, government bonds, real estate, mortgages and mortgage-backed securities, etc. Due to the first point, we are limiting this current study to being focused on the asset allocation problem but will address the second point by including many more asset classes in the analysis. We are also limiting the current study to focus on the market risk faced by life insurance companies rather than the myriad of other risks that could manifest themselves, including insurance, credit, liquidity, operational, group, systemic, and regulatory risks.

To do this, though, we need to specifically model the tail of the joint distribution of assets invested in by life insurance companies. Under the classical Gaussian assumption, this is not necessary, even if you are primarily concerned about tail risk, because the whole distribution, including the tail, is fully explained by the mean and variance. However, several studies have provided evidence that this assumption is not supported. Longin (2005) shows that the tails of daily stock returns are generally inconsistent with a Gaussian assumption. Mandelbrot (1963) and Fama (1965) discuss how stock returns are too "peaked" to be normal, which means they have tails that are too heavy. Other studies support the view that the joint dependence across asset classes is unique in the left tail region. For example, Hong, Tu, and Zhou (2007) and Junior and De Paula Franca (2012) observe the phenomenon that correlations of many major asset classes tend towards one during crisis periods and times of market turmoil.

Longin (2005) reviews how one can model the tails of a distribution using extreme value theory. Typically, two approaches may be used to define the tail itself. The first defines an extreme observation (and, thus, located in the tail) to be one that exceeds some threshold, which is typically set by the researcher. For example, the tail may be comprised of all daily returns which are less than the fifth percentile. This definition leads one to use a Generalized Pareto Distribution model. The second definition is based on local maxima or minima where the tail is composed of all observations that are local maxima or minima (e.g., the worst daily return for each month). This definition leads one to use a Generalized Extreme Value Distribution model. These probability distributions will be described in further detail in Section V.C.

Modeling the joint distribution of several asset classes becomes challenging, though, when we move away from a Gaussian assumption. We will work around this obstacle in two ways. First, we will effectively transform the multivariate problem into a univariate one by modeling the tail of portfolios of life insurance company assets. We will build these portfolios by starting with current industry-wide empirical weights and then systematically adjust them to create new portfolios. For each portfolio, the tail will be modeled and the risk-return tradeoff will be analyzed.

Second, we will take advantage of copula theory to model the joint distribution. Copula theory is, in fact, nearly as old as mean-variance optimization given that the key theory was developed by Sklar (1959). His theorem stated that for a set of random variables with continuous cumulative distribution functions, there exists a special function (called a copula) that transforms the marginal cumulative distribution functions into the joint cumulative distribution function. Standard texts on copula theory include Joe (1997) and Nelsen (2006) that cover much of the underlying mathematics and theory. They also describe the rich variety of bivariate copulas that have been developed. However, the number of multivariate copula functions is rather limited. Work by Joe (1996), Bedford and Cooke (2001, 2002), Kurowicka and Cooke (2006), and others outline a method by which this issue can be addressed. By exploiting the recursive decomposition of multivariate density functions into a product of conditional densities, one can build up to the multivariate copula with a series of bivariate pair copulas, which is called a vine copula. This allows one to make full use of the rich variety of bivariate copula functions while dealing with a multivariate problem.

Our contributions are to focus on expanding the range of asset classes when analyzing the asset allocation problem for a life insurance company and utilize a relatively new technique, vine copulas, where it has so far received minimal attention. Typically, the literature on life insurance company asset allocation gives the company a choice of investing in a single risky asset like stocks or possibly up to two assets such as stocks and a money market fund-like investment. However, these fail to cover much of the actual investing activity of the life insurance industry. In order to address the challenges involved with studying this as a multivariate problem in a non-Gaussian world, we also introduce practical applications of vine copula theory to a life insurance setting.

II. Life Insurance Companies

A. Brief History of Insurance on Lives of Persons

A. 1. Early Forms of Life Insurance

Before diving into the actual analysis, though, we find it important to review the nature of the subject at hand, which is a life insurance company, and of its business. Part of understanding this comes by reviewing briefly how providing insurance on the lives of persons developed over time. Relative to the analytical and probabilistic nature of life insurance today, early insurance contracts covering the lives of persons seems very crude and unsophisticated. It was during Roman times that insurance covering the lives of sailors was in use (Bernstein (1998)). However, it was a very different financial product than today. It was essentially a conditional loan. A sailor who needed funds for a voyage would borrow the funds, and repayment of the loan would only occur if the sailor survived the voyage. Thus, this was premium-free insurance and the death benefit, presumably for the benefit of any widow and/or children, took the form of debt relief rather than a lump sum cash payment like today. Nonetheless, the cash flows of the contract were contingent on the survival or death of the covered life and hence can properly be considered an early example of life insurance. This form of insurance also did not originate with the Romans. The Sumerian Code of Hammurabi includes references to "bottomry," which was a conditional loan essentially the same in structure as the Roman sailor's life insurance policy (Bernstein (1998)).

Another example of early life insurance is the Greco-Roman burial society. According to O'Donnell (1936), burial societies date to the second and third centuries before Christ in Greece, and their purpose was to provide for a decent burial as well as the continuing needs of widows and orphans. These burial societies continued into the Roman era where soldiers, nobility, and even the lower classes could join, contribute to a joint pool of funds, and thus provide for a proper burial. In fact, the Romans believed that the soul of the deceased would find no rest without a proper burial (O'Donnell (1936)).

These early forms of life insurance lacked some key elements that would later facilitate the development of life insurance as a stand-alone business. The providers of life-contingent loans did not have any ability to systematically and accurately price the guarantee being provided, very likely did not create much of a risk pool to diversify the risk of claims, and lacked a broader applicability outside the scenario where an upfront capital investment (hence, the loan) is needed. For example, bottomry would not be useful when the covered life is a young and healthy person who is not about to go on a long and perilous maritime journey and requires no loan of funds. In other words, bottomry applied to the risk of death resulting from a specific event related to the loan rather than the risk of death at some eventual but indefinite future point in time potentially many years out. Although the burial societies did effectively create a risk pool by accepting many members, they could not price the guarantee with much mathematical sophistication. In order for these missing elements in the bottomry and burial society contracts to be incorporated, additional theoretical and statistical advances would need to occur in the measurement and modeling of human lives.

A. 2. Theoretical Contributions to the Development of Insurance[1]

There were several theoretical and statistical advances made in the 1600s and 1700s that greatly contributed towards the ability to systematically and profitably provide insurance on the lives of persons. Prior to this time, a life insurance contract was essentially a gamble made by the provider of insurance. There was a lack of analytical tools and intellectual understanding to approach this problem in any way other than using guesswork. The first of these intellectual breakthroughs was the relationship between the probability of an event occurring and the potential consequences of the event should it occur. This particular insight is attributed to a monk at the Port-Royal monastery in Paris by none other than Blaise Pascal. In the *Ars Cogitandi* (*Logic, or the Art of Thinking*), a short and simple passage makes it clear that "fear of harm ought to be proportional not merely to the gravity of the harm, but also to the probability of the event." Applying this to the case of life insurance, the French monk was saying that the price one would be willing to pay for life insurance is a function of two variables, the probability of unexpected death and the monetary consequences of such a death. Given that a life insurance provider is promising to relieve its policyholder of the "fear of harm" where the harm in this case is untimely death, the value of such a contract to the policyholder reflects the "fear of harm" felt by the person sans insurance. The value of life insurance to a young, healthy, strapping lad with no wife or children is very likely to be significantly smaller than it would be to a middle-aged and sickly man with a wife and several children who would have almost no source of labor income should the man die. As a result, the young, single man would be willing to pay much less for insurance coverage on his life than the older man with many dependents.

Although the intellectual contribution provided by the Port-Royal monk helped specify the appropriate relationship between the probability of death, the consequences of death, and the

[1] The information in this section is credited to Bernstein (1998) and Ferguson (2008) unless otherwise cited.

value of life insurance, there was still the pesky problem that very little was known about the probability of an untimely death. Obviously, an insurance provider even many centuries ago would be able to tell a difference in the probability of death for extreme examples, like the young man mentioned above and an elderly man who is practically on his deathbed. It does not take any statistical insight to know that the young man has a relatively low chance of death and the elderly man has a relatively high chance of death. However, there would still be a significant lack of precision in these estimates. Exactly how wide is the range between relatively high and relatively low chance of death? Supposing that the probability for the elderly man is approaching one, is the probability for the young man almost zero, 0.10, 0.25, or even as high as 0.50? All of those could be considered significantly lower probabilities than 0.99 or one depending on the probability distribution of death. Even so, this provides very little help in distinguishing between the probabilities of various men or women that would likely fall somewhere in the middle of the distribution.

Thus, a life insurance provider for a certain population of people needs to have a somewhat sophisticated understanding of the probability distribution of death for many or all of the members in the population. Some of the seminal work on specifying this population distribution was done by two members of the Royal Society in England, John Graunt and Edmund Halley. In truth, these men were building upon a foundation laid many years earlier by the Roman prefect Domitius Ulpianus, as presented to us by Clendenin (1932) and O'Donnell (1936). For inheritance purposes, Roman authorities needed to value annuities being passed on to heirs. Although the Ulpian life table is the earliest such table known to us, it is unclear how much the figures were based on mathematics and statistics. It "was most probably based on recorded observations of the value of annuities rather than on the number of deaths occurring within a given time at various ages" (Clendenin (1932)). Nonetheless, it would contain the best estimates of mortality for centuries until the two Royal Society members took up the challenge.

Earlier work was provided by Graunt in which he compiled counts of the number of deaths and births and the causes of death in London for 1604 to 1661. To do so, he used the bills of mortality that the city started collecting in 1603. Apparently, this was inspired by some concept of what we would now call market research. Himself a merchant, Graunt noted a benefit of his study being "to know how many People there be of each Sex, State, Age, Religion, Trade, Rank, or Degree, &c. by the knowing whereof, Trade and Government may be made more certaine and Regular; for, if men knew the People, as aforesaid, they might know the Consumption they would make, so as Trade might not be hoped for where it is impossible" (Graunt (1665)).

Graunt's work also has the benefit of being very detailed in some respects, particularly in terms of the causes of death. For example, on at least one occasion, the bills of mortality attribute the cause of one's death to being "Burnt in his Bed by a Candle at St. Giles Cripplegate." There was also a particular interest in detailing the effect of the Black Plague on London. For each week, the bills of mortality listed the number of deaths due to the Plague, the number of Parishes that were clear of the Plague, and the number of Parishes infected by the Plague. In one particularly horrible week, in September 1665, a total of 7,165 people died from the Plague and only four of 130 parishes were clear of it. In contrast, only 344 people died from all causes in April of that year and no parishes were infected by the Plague.

Such a detailed and long account of the deaths of people in London certainly provides a historical record upon which one could base an estimate of the likelihood that a certain number of people will die in the city in any given week or year. However, Graunt attempted to go even farther. By making some key assumptions, Graunt estimates the likelihood of living to some specified age. Based on his statistical work, he produced the probabilities in Table 2.1 (on the following page) that could be used to estimate the likelihood of death for people of various ages.

Graunt's work had its limitations, though. First, he himself admitted that the diagnosis of the cause of death was uncertain at best given the limited ability of medicine at that time. Second, he used the number of baptisms in a given week to estimate the number of births, but he only included the baptisms from the Church of England. Certainly this would capture most of the births in London at this time, but any births occurring among Catholics or others not affiliated with either church would have been excluded. Third, his data did not include the ages at death, so he lacked the evidence to more conclusively determine the probabilities of living to various ages.

It would be a fellow member of the Royal Society, Edmund Halley, who would provide more definitive evidence on life expectancies. Halley, who the famously regular comet is named after, decided to engage in a similar task of chronicling the births and deaths of a particular population in order to better understand the likelihood of death. In order to extend the work of Graunt, Halley chose to study the data from another city that kept better records. He studied the records of the town of Breslaw, now called Wrozlaw in Poland. The records of Breslaw provided significantly more detail on the ages at death, and using this data, Halley was able to provide much more precise estimates of life expectancies. In fact, Halley found that Graunt's estimate of the likelihood of surviving beyond six years from birth (64%) for London to be optimistic for Breslaw. In Breslaw, he found that only about 56% of those born survived at least six years. Halley also realized that his work on the likelihood of death for persons of various ages had a very practical application to life insurance products. He included a discussion of the valuation of annuities in light of his results from the population life expectancies study. Alas, England was not quick to revise their annuity selling practices for quite a while after Halley published his results. It would take nearly a hundred years before the English government stopped selling annuities at the same price to everyone regardless of age. Nonetheless, the work of Halley and Graunt laid the foundation for the use of actuarial analysis to estimate the life expectancy of a person seeking life insurance coverage, and so, it has been crucial to the development of proper and stable pricing of life insurance products. The tables produced by these early forecasters of mortality are re-presented in Table 2.1 below.

Table 2.1. Early Life Tables

This table contains the life expectancies of persons as given by some of the earliest known life tables. John Graunt, in the 17th century calculated the likelihood that a person will survive until certain ages. Domitius Ulpianus, in the 2nd-3rd centuries, estimated the future life expectancy of a person given they have already survived to a certain age.

Graunt		**Ulpianus**	
Age	*Survival Probability (%)*	*Ages*	*Life Expectancy (years)*
0	100	0 – 20	30
6	64	25 – 30	25
16	40	35 – 40	20
26	25	41 – 42	18
36	16	43 – 44	16
46	10	45 – 46	14
56	6	47 – 48	12
66	3	49 – 50	10
76	1	55 – 60	7
		60+	5

A third key theoretical contribution to the development of life insurance was work done by Jacob Bernoulli around the turn of the 18th century. Bernoulli sought to better understand how one could develop estimates of the probability that a certain event occur based on a finite number of samplings. When it comes to a game of chance, such as rolling a die, the theory of probability can exactly measure the chance that a certain event will occur. For example, there is exactly a 1/6 probability that a three will be rolled with a fair die. However, such a precise understanding of its own mortality has not been granted to the human mind. We cannot measure with exactness the probability that a particular person will live until next year or even that 95 out of every 100 policyholders will live to next year. In fact, the famous mathematician Gottfried Wilhelm Leibniz expressed to Bernoulli his skepticism in improving this state of affairs. "[N]ature has established patterns originating in the return of events, but only for the most part. New illnesses flood the human race, so that no matter how many experiments you have done on corpses, you have not thereby imposed a limit on the nature of events so that in the future they could not vary." Bernoulli, though, sought not for absolute certainty but rather sufficient certainty in order to make it useful.

As we discussed before, Graunt and Halley had already done work in estimating the likelihood of survival and death based on historical data. The key contribution provided by Bernoulli was in knowing how certain we can be in making these estimates. After all, there is no guarantee that the true probability of an event will be revealed with any sample of data even for games of chance. Out of any six throws of a die, you may not roll a three even once or you may do so more than once. However, if you throw the die enough times and record how often you roll a three, the likelihood of rolling a three based on the data will start to converge on the true probability of 1/6. At some point, you could be *reasonably* certain that the probability of rolling a three on any given roll is 1/6. True, you could not be absolutely certain based only on the data, but you could be certain enough in your estimate in order to start making decisions based on it. Bernoulli was seeking after what he called "moral certainty" rather than absolute certainty.

His quest for moral certainty led Bernoulli to a result that is called the Law of Large Numbers. This is a tremendous result for the purpose of statistical inference because it guides us in determining how much confidence we can place in a given estimate based on a sample of data. The Law of Large Numbers tells us that it is more likely for an estimate based on a large amount of data than for an estimate based on a small amount of data to differ from the true value by less than some specified margin. In other words, we can place more and more confidence in our estimate as we increase the amount of sampling. It allows us to express our degree of moral certainty in the estimate by placing a confidence interval around it. For example, we estimate that 97% of an insured population will survive the coming year and we are 95% confident that the true probability is somewhere between 96% and 98%.

Bernoulli's contribution is as important as the work of Graunt and Halley. If Graunt and Halley provided us with an ability to estimate the likelihood that a certain number of insured people will die in the coming year, Bernoulli gives the ability to determine whether or not that estimate is worth using. Insurance providers need estimates they can depend upon to make decisions. They cannot long survive if they unwittingly take on too many bad risks at a low price. Avoiding that requires confidently estimating the true probability of death among those they are insuring, as morbid as that may sound, so that they can set a fair price for the risk they are taking on. After all, the risk of untimely death does not vanish when a life insurance policy is purchased but is transferred from one party (the insured) to another (the life insurance provider). As a result of Jacob Bernoulli's work, we can know, with moral certainty at least, that our estimate of these risks is sufficiently precise to be useful for making decisions.

A. 3. Scottish Origins of Modern Life Insurance

Although mathematicians and statisticians like a certain monk of Port-Royal, John Graunt, Edmund Halley, and Jacob Bernoulli provided a theoretical bridge to span the gap from the unsophisticated and scattered provision of life insurance to its modern form, it would take Scottish Presbyterians to actually cross that bridge. Two ministers, Robert Wallace and Alexander Webster, of the Church of Scotland were particularly troubled about the plight of widows and children of those ministers who met a premature death (Ferguson (2008)). Ultimately, they set up the first life insurance fund that resembles the provision of life insurance today. Instead of paying out claims from the annual premiums paid in by the ministers, they decided to build up a fund, invest it, and then pay out claims primarily from the investment returns. In order to properly price the life insurance coverage, they needed to accurately estimate the number of beneficiaries of the insurance in the future and the amount of money needed to support them. Drawing on the earlier theoretical contributions provided by Graunt, Halley, and Bernoulli, these two ministers were able to make these calculations.

The “Fund for a Provision for the Widows and Children of the Ministers of the Church of Scotland” got started in 1748 and the scheme quickly caught on. Similar insurance funds were started outside of Scotland (the Presbyterian Ministers’ Fund of Philadelphia in 1761) and by those in other professions (the United Incorporations of St. Mary’s Chapel in 1768 for Scottish artisans). Perhaps just as importantly for the development of life insurance, the idea of having one’s life insured became ingrained in the culture of thrifty Calvinist Scotland (Ferguson (2008)). As a result, it was shown that size is important to the provision of life insurance. As Bernoulli surmised, it is easier to confidently estimate the claims needing to be paid out in any given year when the size of the insured population is larger. No longer do the providers of life insurance need to be gamblers and testers of fate as before.

A. 4. Ensuing Development of Life Insurance as a Business[2]

Following the emergence of modern life insurance in the mid-18th century, life insurance as a business did not experience sudden and significant growth for some time. This was due to multiple factors including a general apprehension towards life insurance as a concept (it was viewed as trying to assign a monetary value on someone's life), legal restrictions that inhibited the ability of a widow to collect the death benefit on her husband's policy, and insurer policies that sought to mitigate their risk by restricting the activities of the policyholders. This started to change, though, when the idea of forming a mutual insurance company caught on following the Panic of 1837. By setting up a mutual company, the life insurer could market the ability of policyholders to become owners of the business and share in the profits through either higher dividends or reduced premiums. This helped produce much industry growth but insurers started engaging in some fraudulent activities to try to survive the stiff competition. In response, states started regulating the life insurance industry through capital and reserve requirements and some consumer-friendly laws and regulations.

The resulting increase in consumer confidence regarding the stability of the industry and overall economic expansion during the latter part of the 19th century brought a new wave of strong growth in life insurance. Again, this renewed expansion brought with it fresh accusations of mismanagement and fraud. The Armstrong Investigations of 1905 in New York set the stage for new regulations that included banning the ownership of common stock and underwriting securities by life insurers. The restriction on common stock ownership appears to have held until the beginning of Separate Account policies in the 1950s (Hart (1965)). As a result of being restricted from engaging in many investment-type activities, life insurers competed in the early part of the twentieth century by developing new product lines such as group insurance and

[2] The information in this section is credited to NAIC (2013) unless otherwise cited.

annuities and key personnel insurance. Since life insurers could not own common stock, they actually did not get hit as hard as some other financial institutions following the stock market crash at the onset of the Great Depression. Only about 6% of life insurers went into receivership while more than 15% of banks failed, and even the policyholders of the failed insurers had their claims paid in full due to reinsurance agreements while depositors of the failed banks lost about $1.3 billion.

A new era in life insurance started in the mid-1950s when TIAA-CREF issued the first variable annuity products. By making the rates of return earned by policyholders depend on the performance of underlying investments, the variable annuity enabled life insurers to transfer some risk from the company to the policyholder and market an ability to better handle the rising interest rates of that time. Competitive pressure from other financial products during a high interest rate environment also pushed life insurers to develop other products that are still sold today. These include variable and universal life insurance, which were developed in the 1970s and 1980s. By making rates of return more sensitive to movements in interest rates or the equity market with these new products, life insurers could better compete with other financial products such as money market funds, mutual funds, and U.S. Treasury securities.

In recent years, the biggest trends in the life insurance industry appear to be the increasing prominence of investment-type products, such as the variable annuity, and steady demutualization. Demutualization is the process by which an insurance company converts from being a mutual company owned by the policyholders to a stock company owned by stockholders. A number of major life insurers including John Hancock, MetLife, and Prudential of America have undergone this conversion since the beginning of the 21^{st} century. A big reason for the shift towards stock ownership is due to the relative difficulty of a mutual company to raise capital.

They are only allowed to do so through retained earnings or by issuing a particular type of debt called a surplus note (Viswanathan and Cummins (2003)). This is also related to a wave of consolidation within the broader financial services industry. By converting to a stock-owned company, life insurers can also participate in the merger and acquisition activities happening throughout the industry.

B. Overview of the Life Insurance Business Today

B. 1. Life Insurance Products

The products sold by modern life insurance companies can be classified into two broad groups. One of the groups is of course life insurance. They continue to sell insurance on the lives of persons as they always have. The other group includes the various types of annuity products that are now sold in addition to traditional life insurance. In fact, annuities have surpassed life insurance products in terms of premiums received (the source of sales revenue for the life insurance industry). According to the American Council of Life Insurers ("ACLI"), the net premium receipts from annuity products were about $287.7 billion compared to only $130.6 billion for life insurance during 2013 (ACLI (2014)).

Over time, the life insurance industry has continued to develop its life insurance products beyond the basic coverage offered by the Church of Scotland in the 18th century. In fact, the structure of that coverage is essentially nonexistent today as a single product. That product was purely insurance coverage on the lives of any ministers who purchased it, and the coverage lasted until they died regardless of how long that took (i.e., the coverage had a for-life term). Today, the only life insurance products that are purely insurance coverage are term life insurance policies. However, these policies will only provide coverage for a set number of years. The products that have a for-life term now come with some type of savings vehicle embedded in them. These include whole life insurance, universal life insurance, and variable life insurance.

Term life insurance is purest form of life insurance typically available today. However, the coverage comes for limited term, or length of time. The policy's coverage would start at the policy's inception and continue until the earlier of the term's expiration (e.g., 20 years after inception) or the covered person's death. If death occurs prior to the term's expiration, then the corresponding death benefit would go to the beneficiary of the policy. Beneficiaries are determined when the policy start and often include spouses, children, other family members, or even a non-profit organization such as the policyholder's church. Premiums charged on term life insurance are very sensitive to the age of the covered person at inception. For those desiring coverage that are younger and healthier, premiums can be very low indeed since there is a low probability that the death benefit will be paid out. The premiums would progressively rise with age at inception and inversely with health of the covered person. As a result, term life coverage for older and sicker applicants could become quite expensive, again reflecting the likelihood that the death benefit will be paid out. In accord with the monk at Port-Royal many years ago, the value of the insurance coverage significantly increases as the probability of needing that insurance rises.

Whole life insurance provides insurance for the full remaining life of the covered person and allows them to accumulate savings in the policy (the "cash surrender value") as well. Thus, premiums on whole life policies for younger and middle-aged policyholders tend to be much more expensive than on a term life policy with a similar amount of insurance coverage. This is due to two factors. First, a covered person is guaranteed to die while the insurance coverage is in effect with a whole life policy while he/she is not with a term life policy. Second, the premium must be higher than the cost of providing insurance coverage with a whole life policy in order to build up the cash surrender value. The premiums charged on a whole life policy remain constant from year to year although some policies may have the policyholder only pay premiums for a fixed number of years (NAIC (2007)).

Universal life insurance also provides insurance coverage for the full remaining life of the covered person and builds up cash surrender value, but there is more flexibility built in to the premiums charged. With this type of life insurance, the premiums go into an account on which the life insurance company pays interest. From this account, the company would also deduct the periodic cost of insurance and any other charges. As long as there is sufficient money in their account to cover the periodic deductions, policyholders need not pay any additional premiums (NAIC (2007)). Of course, deciding to forgo premium payments will also result in a reduced cash surrender value.

Variable life insurance is a combination (some might call it an unholy alliance) of a life insurance policy and an investment account. As with whole life and universal life insurance, the policyholder pays premiums to the company and builds up cash surrender value. The defining feature of a variable life policy is that the policyholder is able to invest their savings in one or more investment options allowed by the policy (NAIC (2007)). For example, there might be a variety of mutual funds available to choose from as in a 401(k)-style retirement plan. If the chosen investments perform well over time, then the policyholder will benefit by having higher cash surrender value and higher death benefits than with the other types of for-life policies. Of course, that sensitivity extends to the downside as well. If the policyholder makes poor investment choices, they would be left with less cash surrender value as well as a reduced death benefit. To mitigate some of the effects of the poor performance scenario, the policyholder may receive a guaranteed minimum death benefit in exchange for a somewhat higher premium.

If the purpose of a life insurance product is to protect the covered person's beneficiaries against the risk of early death, then annuity products seek to protect the covered person against the risk that death is overly-mature in some sense. Persons in a retirement phase of life often have a significant exposure to the risk that they will live longer than expected and completely draw down their retirement savings prior to death. Annuity products serve to mitigate this risk by paying out a regular stream of income for either the remaining life of the covered person or for a specified number of years. Obviously, the for-life annuity stream provides the greatest protection against the risk of outliving your savings.

Annuities are now sold as one of two main types, as either a fixed annuity or a variable annuity. Both of these types have an accumulation phase and an income phase. During the accumulation phase, the policyholder pays premiums to the insurance company and accumulates a pool of money within the policy. When the accumulation phase ends and the income phase begins, the insurance company would take the full amount of accumulated savings and promise to start making income payments of a certain amount. The expected present value of the income payments, as discounted by some fixed rate of return, would correspond to the savings built up during the accumulation phase. Thus, as the policyholder pays more premiums, the income payments that will be paid out during the income phase rise as well. These income payments would then continue for either a guaranteed number of years as stated in the policy or for the remaining life of the policyholder. A variation on this setup is a joint annuity where the income payments continue for the longer of the remaining lives of the covered person and his/her spouse.

With a fixed annuity, the policyholder pays premiums to the insurance company who then provides fixed rates of interest with a minimum guarantee. During the accumulation phase, this means that the rates of return earned on the accumulated premiums are fixed rates set periodically by the insurance company. The income payments are then determined based on a fixed rate of return corresponding to the general level of interest rates when the income phase begins. When the accumulation phase is drawn out, it is called a deferred fixed annuity because there is a time lag between the premiums and the start of the income payments. Alternatively, when the income payments begin immediately, it is called an immediate fixed annuity.

Another variation on the basic fixed annuity setup is a fixed index annuity and allows for greater rate of return potential while maintaining downside protection. Like a basic fixed annuity, the policyholder pays premiums during the accumulation phase and receives interest on the accumulated policy balance. With a fixed index annuity, though, the interest rate is now tied to a benchmark index such as the S&P 500. This exposure to the equity market provides the policyholder with the potential to earn higher interest rates than would otherwise be available on a fixed annuity product. However, the exposure is of a very limited nature through the use of a floor and a ceiling on the rate earned. For example, if the floor is 3.00% and the ceiling is 5.00%, then the policyholder gains a little bit of upside potential without needing to take on too much downside risk.

A newer, non-traditional type of annuity product is the variable annuity. This is essentially an investment vehicle living within an annuity structure. It could be considered the annuity-version of variable life insurance. Like other annuities, it has an accumulation phase and an income phase. During the accumulation phase, policyholders pay premiums to the insurance company. Instead of being credited with a fixed interest rate or a rate based on the limited performance of a specified index, the policyholder of a variable annuity is allowed to invest their accumulated policy balance in a range of investment options. One option may offer a fixed rate of interest guaranteed by the insurance company for certain term. However, the other options allow the policyholder to invest in a wide variety of asset classes including domestic stocks and bonds, international stocks and bonds, money market accounts, and alternative investments. Thus, a variable annuity policyholder has significantly more upside potential than other annuity products. Variable annuity products often allow policyholders to choose from a range of elective guarantees that help protect them against downside risk as well. These take the form of various types of guaranteed minimum death benefits, guaranteed minimum withdrawal benefits, and guaranteed minimum income benefits. The death benefits would pay out at least a guaranteed minimum amount if the covered person dies prior to the income phase. The withdrawal benefits allow the policyholder to withdraw a guaranteed minimum amount on a periodic basis. In some sense, these give the policyholder an ability to "pre-annuitize," or start receiving regular income payments during the accumulation phase. The income benefits give the policyholder a guaranteed minimum annuity payment during the income phase. The market for variable annuity products has been greatly benefitted by the favorable tax treatment of annuities. Similar to a standard 401(k)-type retirement plan, variable annuity policyholders are able to build up account value on a tax-deferred basis.

There are interesting risk implications of the significant increase in annuity business for life insurance companies. Given a particular applicant for either a life insurance product or an annuity income stream, the product's pricing will depend upon the applicant's expected remaining life. This estimate would likely come from an updated mortality table very similar in spirit to those published by Graunt and Halley over 200 years ago. As with all estimates, though, there is uncertainty and some probability distribution extends to the left and right of the estimated remaining life. With life insurance, the policyholder seeks to protect against the risk that their true remaining life is in the left tail of this distribution (i.e., their actual remaining life is much shorter than they expect). With annuity payouts, the policyholder seeks to protect against the risk that their true remaining life is in the right tail of this distribution (i.e., their actual remaining life is much longer than they expect). The policyholder protects against the corresponding risk by purchasing life insurance and/or annuities, and thereby transfers the risk to the life insurance company. Hence, as a greater amount of business for life insurance companies is coming from annuity sales rather than life insurance sales, they are starting to take on the risk from not only the left tail but from both the left and the right tail. In some sense, they are taking on a short straddle, or short volatility, position regarding the estimated remaining lives of their policyholder population.

B. 2. Scale and Scope of Life Insurance Companies

Life insurance companies are a major component of and investor in the national economy. First of all, the industry's products provide financial stability and security to millions of households. In total, 75 million households in the United States have life insurance and annuity coverage through life insurance companies, which accounts for 66% of all households (Ernst & Young (2014)). On all life insurance and annuity products, the industry paid out about $411.6 billion during 2013 alone (ACLI (2014)). For life insurance alone, the total face amount of in-force life insurance is about $19.7 trillion as of 2013 (ACLI (2014)). This is comparable to the total book value of all domestic corporate bonds and U.S. Treasury securities, which was about $21.2 trillion as of 2013 (Board of Governors (2014), Table L.209).

Life insurance companies also have a wide scope based on their assets and investment activities. In total, life insurance companies held financial assets totaling almost $6.0 trillion at the end of 2013 (Board of Governors (2014), Table L.116). In addition, these investments often fund crucial long-term capital in the economy. Life insurance companies fund about 20% of the corporate and foreign bond market and about 12.5% of the commercial mortgage market as of the end of 2012 (Board of Governors (2014), Tables L.212 and L.220).

At a company level, life insurance companies can operate on a significant scale. Some of the largest American life insurance companies include Prudential Financial, MetLife, Aegon, Jackson National Life Insurance Company, and Lincoln National. These five companies alone account for nearly 28% of the total direct premiums paid to the industry in 2013 (NAIC (2014)), and hold assets in the hundreds of billions of dollars. Although not as large as the biggest banks, the big life insurance companies have attained a scale that places them among the major financial businesses in the country.

B. 3. Ownership of Life Insurance Companies

The ownership of a life insurance company generally takes one of two forms. First, the company could be owned by stockholders who each own shares of stock issued by the company. Life insurers with this type of ownership structure are called stock life insurance companies. This is the same ownership form as publicly traded corporations, and in fact, some of the stock life insurance companies trade on the secondary market exchanges themselves. Of course, the usual corporate governance issues of other stock-owned companies come along with this structure. As a result, stock life insurance companies are exposed to potential agency costs due to conflicts of interest between the shareholders and the managers.

The other form is akin to the ownership structure of credit unions in the banking sector. In this case, the life insurance company is formally owned by the policyholders and no stock is issued to the public. These life insurers are called mutual life insurance companies. An argument for governing a life insurance company with customer-owners is that the insurer would operate for the benefit of the policyholders. Since the policyholders face potentially extreme financial dislocations if the insurer fails, this could be an important consideration. However, mutual life insurance companies are also exposed to potential agency costs, as with the stock life insurers, but between managers and policyholders instead.

Both forms of ownership structure make up significant portions of the life insurance industry. Stock-owned companies comprise 76% of all life insurers in the United States and 70% of the total life insurance in force as of 2013 (ACLI (2014)). Mutual companies make up 13% of the industry and have 27% of the total life insurance in force (ACLI (2014)). These relative market weights have not remained stable over time. The trend since the middle of the last century has definitely been in favor of stock ownership. In 1947, mutual companies comprised 69% of the life insurance industry, but their share fell to 43% by 1983 and, as mentioned recently, is now at 14% (Hansmann (1985)).

There are multiple arguments related to why particular companies favor one form of ownership over another. Viswanathan and Cummins (2003) provide a good overview of the various theories and hypotheses. They point out that a clear advantage of the stock-owned form is access to capital. Mutual companies must generally rely on retained earnings to fund new projects and investments whereas stock companies have access to the deep equity and debt markets. They argue, then, that the recent shift into stock-owned companies is due, in large part, to seeking better access to capital.

Stock-owned companies potentially face a conflict of interest not only between managers and shareholders but also between shareholders and policyholders. Hence, taking on a mutual form of ownership may mitigate the conflicts of interest between the shareholders and policyholders since they are one and the same. As Jeng, Lai, and McNamara (2007) mention, though, this may mitigate the shareholder-policyholder conflicts of interest but exacerbate the owner-manager agency costs. They hold that policyholders are less effective monitors of managerial decision making then shareholders. Viswanathan and Cummins (2003) refers to this as the expense preference hypothesis as in mutual companies have less effective monitoring mechanisms for controlling the expense preferences of managers.

In the managerial discretion hypothesis, it is argued that the stock-owned form is better for business activities that require managers to exercise greater discretion in making decisions (Viswanathan and Cummins (2003)). Again, this is due to stock companies having more effective means of mitigating managerial opportunism. Under the assumption that riskier business activities and cash flows with greater uncertainty require more managerial discretion, an implication of this hypothesis is that stock companies would tend to enter riskier lines of business and geographic areas than mutual companies. Lamm-Tennant and Starks (1993) test this and find that stock companies are in fact riskier. Based on the variance of the loss ratio, they find that stock companies have higher total risk but also that stock companies do relatively more business in riskier lines of business and geographic areas.

The maturity hypothesis argues that the mutual form fits better when the insurance activities have a longer expected duration or contract period. The reasoning is that owners and managers may have more opportunities to effectively extract rents from policyholders under longer-term contracts by increasing the riskiness of their asset management, increasing leverage, or otherwise taking on more risk (Viswanathan and Cummins (2003)). By forming a mutual company and making the policyholders the owners, these potential conflicts of interest are avoided.

The final hypotheses relates to the fact that the policyholders in a mutual company, as the owners, have both a fixed claim on the company's assets (through their insurance or annuity policy) as well as the residual claim. In the informational hypothesis, policyholders self-select into the different forms of ownership based on their risk (Viswanathan and Cummins (2003)). Low risks self-select into the mutual companies (since they retain a claim on the remaining surplus after benefits are paid) while high risks self-select into the stock companies. In the risk hypothesis, separating the fixed and residual claimants incentivizes the company to engage in riskier activities. Also, taking on a stock form may also incentivize riskier activities since stock companies can raise needed capital easier and quicker than a mutual company (Viswanathan and Cummins (2003)). The results from Lamm-Tennant and Starks (1993) provide some support for this hypothesis as well.

B. 4. Globalization of the Life Insurance Industry

As with many industries, the life insurance business has been impacted by economic globalization over the past few decades. This is true operationally and competitively as more insurers compete in foreign geographic areas, but it is also the case with regulatory activities. Financial and insurance regulation has become more aligned across national borders through international accords, greater integration of the European Union, and the creation of bodies such as the International Association of Insurance Supervisors ("IAIS").

It could be that greater unification and cooperation of insurance regulation has encouraged the competitive globalization by removing some of the cost of entering a new market. When each country developed its own regulatory framework more or less in isolation from other jurisdictions, then a multi-national insurer must become knowledgeable in multiple, potentially conflicting, regulatory environments. As insurance regulation is globalized, though, the regulatory environments become more uniform, particularly across the developed economic markets, and thus easier to enter by foreign firms.

Whether or not increasing competition from foreign insurers is the result of regulatory globalization, it is apparent that domestic insurers in many countries must compete against significantly more foreign firms now. Figure 2.1 charts the weighted average of foreign insurer market share of 31 countries[3] of the Organisation for Economic Co-operation and Development ("OECD") with the weights based on each country's market share based on premiums. Because data are not provided for every country-year observation over the chart's time period (1989 – 2012), the weighted average for each year includes only those countries with both foreign market

[3] The countries include Australia, Austria, Canada, Chile, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Japan, Korea, Luxembourg, Mexico, the Netherlands, Norway, Poland, Portugal, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, the United Kingdom, and the United States.

share and country premium market share data for that year. Weights are determined in each year using only these countries as well.

Figure 2.1. Foreign Life Insurer Market Share in OECD Countries
This graph plots the weighted market share average of foreign life insurers from 31 OECD countries from 1989 through 2012. The weights used in the average are based on each country's global market share of premiums. For each year, the average includes only those countries for which both foreign insurer market share and country-level premium data are available. Data are from OECD (1998, 2006, and 2013).

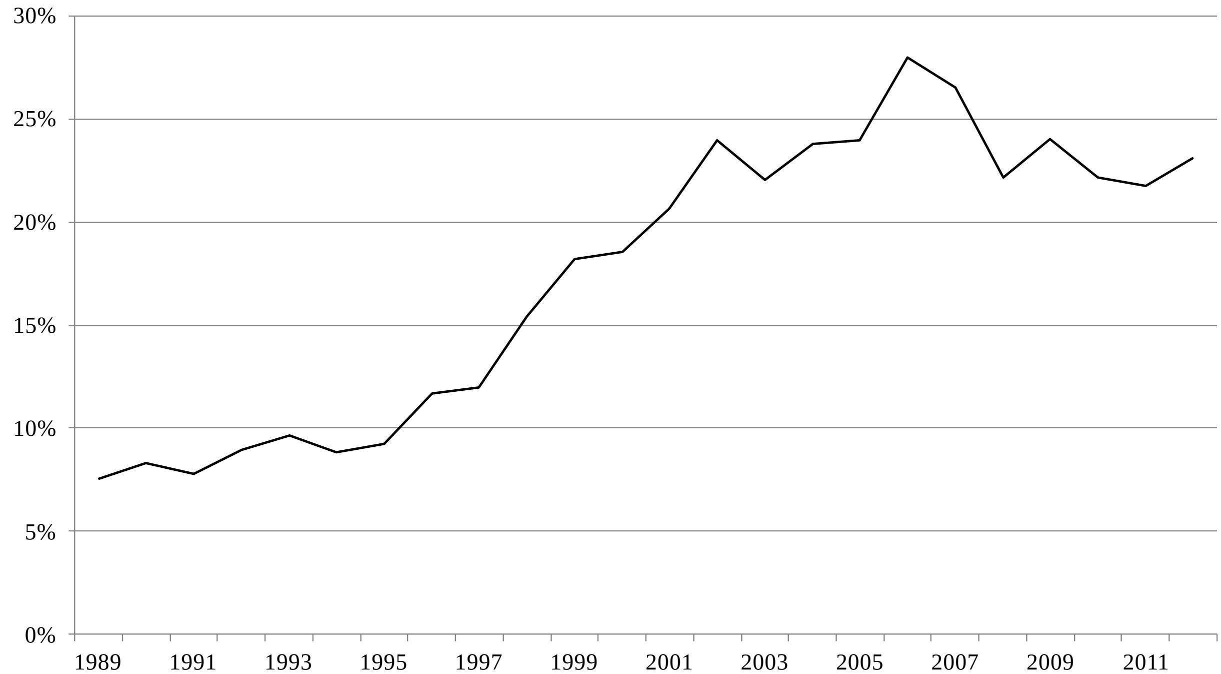

From this chart it is clear that the average market share captured by foreign life insurers has dramatically increased over the past 20 to 25 years. Moreover, this shift occurred entirely over roughly a ten-year period starting in 1995. In that year, the average foreign market share was 9.24%, and it had risen to a high of nearly 28% in 2006. Since then, it has receded a bit from that height but has largely sustained the higher level.

What is the significance of this globalization of the life insurance industry? For one thing, it should increase the competitiveness of each domestic market experiencing this shift as more insurers enter. Since these are multi-national firms, they likely have ample resources and a solid capital base, which makes them relatively tough competition. It may help diversify the foreign insurer's insurance risks by entering a new geographic area. The benefit is more limited than for non-life insurance, though, because all people everywhere will die and not all cars or houses will have insurance claims. Still, it may help the insurer to move into an area with higher (for life insurance) or lower (for annuities) life expectancies than their current market. However, there are challenges for the foreign insurer too. It may have to establish name recognition and, especially for insurers who focus on newer variable policies, manage differing levels of policyholder risk aversion than in their domestic market. Thus, one strategy for a foreign insurer to enter a new region is to purchase a company already located in and operating in the new market.

B. 5. Assets and Liabilities of Life Insurance Companies

The assets of a life insurance company are primarily derived from the premiums paid in by policyholders and earnings on prior investments. Given that the life insurance business revolves around the selling of financial contracts, only a small fraction of the company's total assets consists of tangible, physical assets such as buildings, equipment, and supplies. Most of the life insurer's assets are invested in financial securities such as stocks, bonds, and mortgage securities. As mentioned earlier, the scale and scope of the life insurance industry combined with the fact that most of its assets are invested in financial securities allows it to be a major participant in these markets and provide significant long-term capital to other sectors of the economy.

These financial assets ultimately support the amounts that will be paid out to policyholders or their beneficiaries for life insurance death benefits, annuity payouts, etc. Thus, the assets of a life insurer are classified into one of two accounts based on the nature of the contractual obligations they support (ACLI (2014)). The General Account assets support the payouts on fixed-payment products such as life insurance or a fixed annuity. The Separate Account assets support the payouts on products associated with policyholder investment risk such as variable life insurance and variable annuities. Because the policyholders of these variable products are given the ability to allocate their premiums among several investment choices, the asset allocation of the Separate Account reflects the aggregate investment choices of those policyholders whereas the asset allocation of the General Account reflects the investment decisions of the company.

It is very interesting to compare the two sets of asset allocation decisions made in the General and Separate Accounts. The General Account asset allocation reflects traditional life insurance company investment policies, which means it is heavily weighted bonds with much of the remainder invested riskier securities (mainly stocks and mortgages) to provide some additional upside potential. In total, there was about $3.8 trillion in the General Accounts of the life insurance companies as of 2013 (ACLI (2014)). Of this, nearly 71.0% is invested in long-term bonds (including mortgage-backed securities) while direct mortgage investments make up about 9.6% and stocks receive a paltry 2.2%. In contrast, only 12.9% of the nearly $2.35 trillion in Separate Account assets is invested in long-term bonds and 81.7% is invested in stocks.

Table 2.2. Aggregate Life Insurer Balance Sheet Composition

This table shows the aggregate balance sheet asset distribution of the life insurance industry and broken down into the "General Account" (assets directly under the control of the company) and the "Separate Account" (assets for which policyholders direct investment through certain insurance and annuity policies). It shows the distribution of assets for 1999, 2008, and 2013. The Other Assets category includes short-term investments, cash and cash equivalents, derivatives securities, and accounting assets such as premiums owed and interest earned but not yet received. The 1999 and 2008 data are from ACLI (2010) while the 2013 data are from ACLI (2014).

General Account			
	1999	*2008*	*2013*
Bonds	71.21%	67.76%	70.97%
Stocks	5.08%	3.60%	2.23%
Mortgages	11.37%	10.34%	9.56%
Real Estate	1.28%	0.62%	0.60%
Policy Loans	4.94%	3.73%	3.46%
Other Assets	6.12%	13.96%	13.18%
Separate Account			
	1999	*2008*	*2013*
Bonds	13.28%	15.47%	12.89%
Stocks	81.15%	73.90%	81.67%
Mortgages	0.47%	1.06%	0.44%
Real Estate	1.18%	0.88%	0.37%
Policy Loans	0.11%	0.04%	0.02%
Other Assets	3.81%	8.64%	4.62%

These data are shown in Table 2.2 for 2013 as well as corresponding measurements for 1999 and 2008 to compare the recent allocations with those from a stock market peak and trough. Generally, insurers have been reducing their exposures to certain classes of riskier assets like stocks and mortgages since the turn of the century. Instead, they have largely re-allocated funds to the "Other Assets" category, which includes short-term investments, cash holdings, derivatives securities, and accounting assets such as premiums owed to the company and interest earned but not yet received. Policyholders appear to have shifted their equity allocations during the financial crisis, the lower allocation in 2008 could be due at least in part to reduced equity levels rather than actual transfers of funds to other asset classes. As of 2013, policyholders have reverted back to higher equity allocations while shifting away from bonds and other assets, which, again, includes cash and short-term investments.

Figure 2.2. Historical Life Insurer Asset Allocation to Bonds, Stocks, and Mortgages
This graph plots the historical allocations to three primary asset classes for life insurance companies. These asset classes are long-term bonds (including both public and private bonds), stocks (including both common and preferred), and mortgages. It begins in 1917 and continues through 2012, but the data are only provided in five-year increments from 1920 to 1980. To keep the time scale consistent throughout the chart, missing observations from 1920 to 1980 were linearly interpolated using the two closest actual observations. Note that these allocations will not sum to 1 due to excluding other assets. Data are from ACLI (2014).

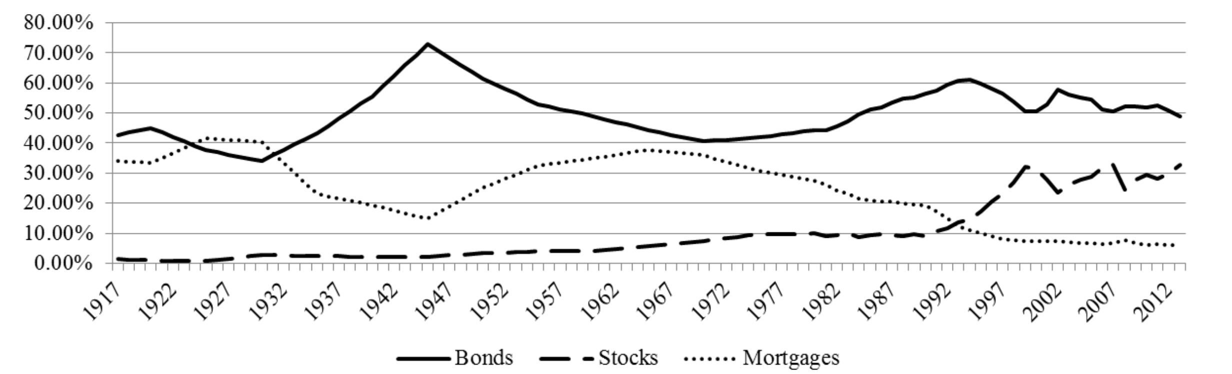

When combining the General and Separate Account assets, we also see an interesting long-term shift in the asset allocation of life insurer assets over the past century (shown in Figure 2.2). In 1917, life insurers primarily invested in two asset classes: bonds and mortgages. In fact, both of these could be thought of as being part of the same asset class (fixed income) depending on how you classify the investable universe. Bonds and mortgages received weights of about 43% and 34%, respectively, while less than 2% of all assets were allocated to stocks. Considering that the traditional life insurance business involves paying out a fairly regular percentage of death benefits each year, focusing on fixed income investments makes sense from an asset and liability cash flow matching perspective.

Following a significant shift into and then back out of bonds around World War II (apparently, life insurers provided some much needed capital for the war effort), these allocations remained roughly the same in 1965 except life insurers had increased their stock proportion somewhat to nearly 6% (largely coming from a reduction in policy loans, which is not shown on Figure 2.2). From 1965, we see two successive long-term shifts in the asset allocation of life insurers. The first, from 1965 to about 1990, was largely a shift from mortgages into bonds. The mortgage allocation went from nearly 38% in 1965 to just less than 20% in 1990. Most of this reduction in mortgage investment shifted into bonds, but there was an initial increase in stock allocations at the beginning of this period. The cause for this shift is unclear. Perhaps, life insurers pulled back from financing more of the mortgage market as the government sponsored enterprises (the Federal National Mortgage Association and Federal Home Loan Mortgage Corporation) entered. Interestingly, it is around 1965 that the data starts including Separate Account assets, so it may also reflect the fact that policyholders typically invest very little of their Separate Accounts in mortgages. The second shift, from 1990 to 2000, was a large increase in equity exposure. Stocks went from a relatively small allocation of about 9% in 1990 to over 30% in 2000. Given that the General Account assets were still mostly invested in bonds and mortgages in 1999 (see Table 2.2), this shift must be largely due to policyholder investments. No doubt policyholders wanted to reap the benefits of the booming stock market over this same time period. The increase in stock allocation largely came from another large reduction in mortgage investments from about 19% in 1990 to about 7% in 2000, although the allocation to bonds also decreased somewhat too. Since then, the allocations for these three major life insurer asset classes have largely been in a holding pattern with some fluctuation around the two market crashes in 2001 and 2008.

Life insurance companies invest primarily in investment-grade long-term bonds at least partially for regulatory reasons. Because the General Account assets support guaranteed payouts to policyholders, life insurers are restricted in how they can invest those assets to minimize the risk of life insurance company failure, in which case the state ultimately becomes responsible for the life insurer's benefit payments up to certain levels. However, there are rational, risk-based reasons for the life insurance companies to continue investing the General Account assets in this manner. It is not uncommon for the policyholders with Separate Account assets to elect certain guarantees on their variable products. Thus, the life insurance company needs to make up the difference between the payouts that can be supported by the policyholder's account and the guaranteed levels. Since the policyholders show an affinity for taking on significant equity risk, which is perfectly rational in the presence of guarantees, it is prudent and rational for the life insurance companies to minimize the equity risk in the assets directly under their control. Otherwise, the company would be inviting a disaster to occur should a significant shock in the equity market hit both the Separate and General Account assets at the same time.

The liabilities of a life insurance company primarily consist of monies held in reserve to fund future policy payouts (ACLI (2014)). In fact, the distinction between assets and liabilities here is somewhat abstract as the monies invested in stocks, bonds, and other assets are the same monies held in reserve for future payouts. Life insurers are legally required to maintain reserves at certain levels to provide a reasonable assurance that sufficient funds will be available to make all of the promised payouts. The required levels of reserves are actuarially determined based on forecasts of future premiums, investment gains, policyholder behavior, mortality, and other factors. Of approximately $5.8 trillion in total liabilities for the life insurance industry, about 88.5% consist of these policy reserves.

The remaining liabilities are mostly comprised of other types of reserves. These include reserves for deposit-like contracts, asset fluctuation reserves, and interest maintenance reserves (ACLI (2014)). Deposit-like contracts are those where the payouts are not contingent on the life, death, or disability of the covered person. For example, some annuities promise to pay a certain income benefit every period for a specified number of periods. As a result, the payout stream of cash flows is certain, fixed, and not dependent on the death of the covered person. If the covered person dies unexpectedly prior to the completion of the promised payments, the remaining scheduled payments would be paid to the beneficiary instead. As of 2013, the reserves in place supporting deposit-like contracts totaled $450.4 billion, which is 7.8% of total industry liabilities (ACLI (2014)).

In addition to the required policy reserves discussed earlier, life insurance companies are also required to set aside a certain amount of reserves specifically for investment gains and losses. These are asset fluctuation reserves and interest maintenance reserves. Asset fluctuation reserves cover potential realized and unrealized losses due to defaults on credit securities and equity movements. Interest maintenance reserves cover realized and interest-related gains and losses on credit securities. As of 2013, the industry's asset fluctuation reserves totaled $48.4 billion and interest maintenance reserves totaled $26.5 billion, which are 0.8% and 0.5% of total liabilities respectively (ACLI (2014)).

C. Key Risks Faced by Life Insurance Companies[4]

C. 1. Insurance Risk

One of the primary risks faced by a life insurance company relates to its key operational activity. Insurance risk is the risk that the company's underwriting and/or claims policies and procedures undermine the company's profits, capital position, or reserves. Since selling and pricing insurance, which is the company's underwriting activity, and paying insurance claims are essential to the life insurance business, insurance risk is inherent in the life insurance business to the deepest levels.

Within this risk class, underwriting risk is one of the main components. Underwriting risk is primarily due to the fact that the life insurance company must estimate the risk of and the appropriate price charged to an insurance applicant. First of all, insurance companies face a potential adverse selection problem in that the applicant knows their health and the riskiness of their behavior than the company does. This could lead to a self-selection problem where the people most likely to purchase life insurance are those that have an above-average probability of either premature death (in the case of life insurance) or greater longevity (in the case of annuities).

Even without the adverse selection problem, the life insurance company still faces the risk that their pricing mechanisms do not adequately estimate the inherent risk of selling a life insurance or annuity product to a particular applicant. Without realizing it, the life insurer may be insuring a riskier pool than they priced for. Even creating a less risky pool than they expected could be an issue if other life insurers do not make the same error. In that case, they may be charging too high of a price and losing some business to competitors who are using more accurate pricing.

[4] Much of the information in this section is credited to IAIS (2003) unless otherwise cited.

Another source of insurance risk is when a life insurance company expands into a new geographic area or introduces a new product. Without any past experience to base their underwriting on, the insurer is exposed to the risk that they are not correctly pricing the new business due to greater uncertainty about the inherent risk profile of the applicants. This could be a potentially important risk exposure if an insurer is trying to enter a new market that already has incumbent firms with prior experience. One barrier to entry in a new life insurance market, then, is greater uncertainty surrounding a company's pricing of the new business compared to competitors who are already entrenched in that market.

Related to the risk that the company is not correctly pricing its products is the risk that payouts on their products ultimately deviate from the initial expectations. This is not the case where the insurer misestimated the expected life expectancy or risk of the applicant. Rather, this is the case where the actual payouts differ from the expected amounts, even if the company accurately estimated the expected amounts given the information available at the policy's inception. A special case of this claims risk may be called catastrophic risk, and it is the risk that some catastrophe, whether it is a natural disaster, terrorist attack, etc., causes an unexpected spike in claims. This catastrophic risk would also be an element of the broader tail risk faced by a life insurer.

C. 2. Market Risk

Given the significant scale of the life insurance industry and the amounts of assets and liabilities involved in this line of business, market risk is another very important exposure for a typical life insurance company. Market risk is the risk that movements in the market prices of assets, interest rates, foreign exchange rates, etc. have a negative impact on the capital position and reserves of the life insurer.

Suppose there is a shock in the equity market and stock prices suddenly drop by a fairly significant amount. As a result, the value of the assets in the insurer's Separate Account, which tend to be invested primarily in stocks, takes a significant hit, but the company is still liable to make sure the corresponding guarantees are met. In fact, the value of the liability has risen because the expected amount that will need to be paid by the company out of its General Account has increased. Thus, an adverse market movement such as this has the effect of reducing the total assets, increasing the total liabilities, and eating away at the capital base. In other cases, the effect on the company's capital may be mixed. For example, if interest rates suddenly drop, then the present value of the guarantees goes up (due to a lower discount rate) but the value of the bonds in the company's General Account should also rise.

Given the seemingly constant movements in market valuations for stocks, bonds, and money, this is a risk exposure that life insurance companies need to be monitoring on a regular basis. Even gradual changes that become a significant market movement over time could pose a problem for the financial position of a life insurer if they are not adjusting accordingly along the way.

C. 3. Credit Risk

Life insurance companies conduct business with a number of counterparties that agree to fulfill certain obligations with the company just as the company itself agrees to meet obligations to its policyholders. These counterparties include entities such as companies and governments that promise to make bond payments to the life insurer, reinsurance companies who promise to cover certain claims made on the life insurer, derivative counterparties who agree to take the other side of a derivative transaction, and even policyholders who take out loans on their policies and agree to repay them with interest. Credit risk is the risk that any of these counterparties fail to honor their side of the agreement with the life insurance company.

Credit risk can lead to potentially devastating outcomes for the life insurer. Traditionally, a big source of credit risk is the probability of default on the bonds owned in the insurer's General Account. If these debtors fail to repay the bonds, then the capital and reserves of the life insurer are reduced. A well-diversified bond portfolio can manage a handful of defaults on individual bonds without too much trouble in any given period. If the risk of default increases in a systematic way, such as during a deep recession, then the life insurer's capital position could deteriorate significantly if many defaults occur at approximately the same time. This is a reason that life insurance companies are restricted in the types of bonds they may own. For example, an insurance company would be restricted in its ability to invest in speculative bonds that are rated below investment-grade levels.

Today, credit risk can have a significant impact due to the much greater use of derivative securities over the past few decades. A life insurance company might purchase certain derivative securities to hedge the impact of adverse market movements. Consider again the earlier example where a down shock in the equity market leads to reduced Separate Account values and capital.

If the company had purchased put options, which provide a benefit to the owner as the value of the underlying asset declines, then the company could depend on a cash inflow from the put options to offset the rise in expected liabilities. However, if the counterparty or counterparties of those put options fail to meet the obligations of those contracts, then the insurer could be left exposed to the full impact of the equity shock in addition to losing the premiums paid for the put options. Depending on the size of the market shock, this manifestation of credit risk could have potentially ruinous consequences.

Life insurers can hedge credit risk by being discerning and cautious in the choice of counterparties to deal with. For example, the insurer could choose to only invest in bonds with a credit rating of A or higher. In addition, they can hedge credit risk by being well-diversified in their choice of counterparties. An example of this would be limiting the asset allocation in any one security to be no more than some level such as 5% or by buying derivatives from several counterparties instead of only one.

C. 4. Liquidity Risk

The assets and liabilities of a life insurance company involve many cash flows both coming into the company and going out of the company. A potential source of problems for the company is that these cash flows and additional premiums are not necessarily harmonized together. Liquidity risk is uncertainty related to the timing of the cash flows and the possibility that the company may not have sufficient cash on hand to meet the required policy payments when they need to be paid. It is not an issue of being technically insolvent where the value of the assets is less than the value of the liabilities. The insurer may have ample assets to cover all of their expected liabilities. Rather, it is a problem of having the ability to pay cash out to policyholders on time.

Another meaning of the term liquidity risk is related to this inability to cover payments to policyholders. The alternative type of liquidity risk is the risk that one cannot sell assets quickly except at a steep discount. This can be related to the prior meaning because if a company finds itself in a situation where significantly more claims are being made at a particular time than expected, the company may need to liquidate assets to cover the payments. Thus, it could lead to somewhat of a fire sale situation where unloading a lot of assets on a market at good price could be difficult.

As a result, life insurers may sensibly make asset allocation decisions while monitoring any potential cash flow timing mismatches between the assets and the liabilities. For example, the company may not want to invest only in long-term bonds but diversify across a range of maturities. In addition, the company may want to avoid investing a significant amount of assets in relatively illiquid markets, such as small capitalization stocks in frontier markets, even if they offer a great risk-reward profile. Precautionary actions such as these enable a life insurance company to minimize the probability of finding itself in a tight situation of being asset-rich but cash-poor.

C. 5. Operational Risk

Although the inherent nature of the life insurance business involves promising to pay out certain payments in the future under certain states of the world in return for receiving premiums from policyholders, putting this business into practice necessitates the use of certain internal systems, procedures, and labor services of employees. This creates the potential for these aspects of a life insurer's operations to fail and negatively impact the company's financial position. This is referred to as operational risk.

Operational risk can become manifest in a number of ways. Certain employees could engage in activities such as embezzling company funds, over-promising to potential policyholders, or failing to follow underwriting and other policies. The computerized technology that many of the modern administrative and processing systems rely upon could be undermined by power outages, technological failure, or cyber-attacks. The company's business continuity or disaster recovery plans could be found wanting or inadequate when a triggering event actually occurs. If the life insurer outsources any aspects of the business, then those third party providers, who the life insurer may not have as much control over, may fail to follow the obligations laid out in the outsourcing agreement.

The financial consequences of operational risk can be significant. There may be direct impacts resulting from financial losses due to the actual manifestation of operational risk. For example, if the company's disaster recovery plans are inadequate, then the insurer may lose several days or weeks of normal business operations, revenues, and profits as a result. The financial losses may also stem from any regulatory or legal actions that occur in response to the operational failure. This could be particularly relevant when the operational failure has an adverse impact on all or a group of policyholders.

These consequences are also potentially long-lasting because operational failures can damage the life insurer's reputation and brand. Like many financial institutions, life insurers seek to create a perception that they are financially strong, stable, and prudent. An operational failure can easily undermine such a perception from the public's viewpoint, and it can take years for a company to rebuild their reputation. Thus, operational risk could have a financial impact far beyond the direct financial losses caused by the action itself.

C. 6. Group Risk

Similar to the rest of the broader financial services industry, there has been a fairly significant amount of consolidation activity within the life insurance industry over the past few decades. As a result, many life insurers now must operate within a group setting. This is to say that either the life insurer is owned by a parent holding company that also owns other divisions or lines of business or the life insurer has acquired a one or more subsidiaries that may or may not be engaged in the life insurance business. To be sure, this arrangement can certainly be a source of strength to the life insurer. When the life insurance business is suffering or even finds itself in a financial crisis, it may have access to cheap emergency capital through the parent company. If the life insurer is the parent company, then it may still benefit when it is struggling and non-insurance subsidiaries have strong performance and support the financial results of the overall group. However, a group setting poses risks for the life insurer as well.

The potential for group support can provide a benefit to a life insurer, but it can also be a source of danger. Although the parent company is generally understood to stand ready to support any of the subsidiary businesses in case of need, there may be some discretion involved as well. It could be the case that the parent company is unwilling to support a struggling subsidiary when it finds itself in trouble. Instead, the parent company may prefer to let that subsidiary fail and retain more capital and resources for the remaining lines of business. Also, a group setting increases the risk of contagion as the troubles of one member of the group start to "infect" the other members. One of the life insurer's non-insurance subsidiaries could start underperforming and divert resources away from the insurance businesses.

Joining a group of businesses exposes a life insurance company to the risk of relinquishing some control over its business. The parent company ultimately sets many policies and manages resource allocation across the whole group. Thus, resources and capital may end up being diverted away from the life insurer and toward other members of the group. Management of the life insurer will likely not have full control over setting its own strategies and policies. Instead, management may be constrained by group-wide policies or even somewhat distracted by attending to group initiatives established by the parent company. When the life insurer operates in another jurisdiction from the rest of the group or the parent company, the life insurance company's operations may be affected by the regulatory framework of the rest of the group.

C. 7. Systemic Risk

Like many industries, issues that become evident in one area have the potential to reverberate throughout the whole industry. Although a particular life insurer may not have any responsibility for the original problems, the consequences could have a very adverse effect on its own operations and financial position. Systemic risk refers to these potential spillover effects that are due to a company being one element in a broader system. This risk exposure was particularly acute during the financial crisis of 2007 – 2008. Whole swaths of the financial services industry were avoided for some time simply due to the fact that they were a financial institution and fears about financial institutions were very high.

For life insurance companies, systemic risk could be felt in a few different ways. If a sufficiently significant life insurer falls into financial straits or even fails, that could have an impact on the broader industry. Potential customers may decide to avoid otherwise perfectly healthy insurance companies out of fear that a similar fate may befall them too. Or, the whole industry may start to receive extra regulatory attention as a result of trouble in some areas of the industry. Given that life insurance companies are also included in the broader financial services industry, there could be spillover effects if other financial institutions, such as banks in 2007 – 2008, have significant problems.

C. 8. Regulatory Risk

Like all private enterprises, life insurance companies are under the supervision and monitoring of a number of regulatory agencies and governing bodies. At times, the regulatory authorities may pass down new statutes or rulings that pose unexpected challenges to some or all of the industry's members. The costs associated with these regulatory challenges comprise what we will call regulatory risk.

Although the rule of law, rather than arbitrary bureaucratic power, is important for effective governing in a free society, regulatory actions sometimes produce unexpected costs on private businesses that can upend prior strategies and financial forecasts. A regulator may interpret a legal statute differently than the company expected. New laws, solvency rules, or regulations may be passed that will require life insurers to adjust their plans for growth or investment in new products and lines of business. These are only a couple of examples but there are many ways in which regulation and legal issues can have negative effects on life insurers.

The trend over the history of the life insurance industry has been for regulatory activity to increase with each new crisis leading to a bevy of new agencies, laws, and policies seeking to protect policyholders or maintain stable insurance companies. A recent example of this is the threat of being classified a systemically important financial institution ("SIFI"), which brings with it enhanced monitoring and regulatory attention. Naturally, the response of the industry, like many other industries, is to increase its lobbying activities in an attempt to influence the shape of the regulations it will operate under. Although this opens up the possibility of corruption, there are some rational reasons for the industry to have this influence. Legislators and regulators may not always have the intimate details of the industry or of the insurance companies as the industry itself does, especially given the complexity of many large financial institutions today. If a proposed law or regulation is going to have a particularly damaging impact on the industry or is misguided in its application, then it is reasonable for the industry to seek to influence it toward a better or at least a more palatable alternative.

D. Dependency among the Key Risks

All of these key risks (insurance, market, credit, liquidity, operational, group, systemic, and regulatory) are important sources of uncertainty and potential financial loss for a life insurance company on their own. However, there are interactions and connections between these key risks such that they are not completely independent of each other. The manifestation of one type of risk could lead to consequences of another risk class.

Insurance risks are certainly related to operational and liquidity risks. An operational failure to meet company-wide underwriting standards could result in a riskier pool of policyholders than expected or in larger claims than expected. Or, a technical glitch in the company's pricing software could result in a whole group of policies being mispriced for the true risk being covered. Larger-than-expected claims could lead to a liquidity crisis at the life insurer if the spike in claims is sufficiently severe.

There are also important relationships between market, credit, and liquidity risks. A shock to the credit risk of a company's bond portfolio could also lead to a corresponding shock in the value of the life insurer's bonds and/or stocks. A shock to interest rates could result in larger-than-expected defaults on mortgages similar to the financial crisis of 2007 – 2008. A severe shock to the equity or bond markets could undermine the ability of derivative counterparties to satisfy their contractual obligations with the life insurer. Ultimately, credit and market shocks could result in a sharp reduction in the insurer's capital and possibly lead to a liquidity crisis.

Insurance risks could also be related to group and systemic risks. If the company's underwriting and pricing policies turn out to be particularly poor, a parent company may decide that the issues are deeper than can be fixed by simply providing some additional capital. In turn, the parent company may decide to wind down or divest the insurer as a result. Certain underwriting or pricing practices may become established across an industry such that most or all insurers adopt the industry standard practice. If it turns out that the industry standard was flawed in some key respects, this could be both a systemic and an insurance risk for the insurer.

Regulatory risks could be related to many of these key risks too. If an individual company pursues overly aggressive investment policies or has shoddy underwriting, then it could result in unwanted regulatory action and extra scrutiny. If these risks in become severe and prevalent across the industry, then regulators and/or legislators may develop new solvency standards or otherwise constrain the industry's activities in an attempt to protect policyholders and the industry from current or future crisis.

Undoubtedly, these are not the only connections between these key risks. However, these examples provide insight into how these key risks cannot be treated in isolation from each other. When managing the risks being acquired as a result of insuring the lives of persons, a prudent insurer must consider the potential ramifications that one risk has on other types of risk and the relationships that exist between them.

E. The Ultimate Risk – Individual Life Insurance Company Failure

These key risks can and do create uncertainty and financial problems for life insurers. Typically, manifestations of these risks produce slow growth, reduced returns to owners, or extra regulatory scrutiny. They do not frequently result in the ultimate risk of the life insurance business, and that is the risk that an individual life insurance company fails. Very few interested parties benefit when a life insurer collapses. Shareholders will likely sustain significant financial losses, creditors may not receive full repayment on debt, policyholders may not be able to receive the full payments promised to them, and managers may lose their jobs. However, company failure has happened in the past and will likely happen again.

Just as each state has an insurance commission or department that regulates the life insurers, the states have also established guaranty associations that help protect policyholders should an insurance company become insolvent. In 1983, a voluntary organization called the National Organization of Life and Health Insurance Guaranty Associations ("NOLHGA") was created by the various state-level associations to assist them with multi-state insolvencies (NOLHGA (2014)). Similar to how the Federal Deposit Insurance Corporation guarantees the account values of bank depositors up to a certain level, the guaranty associations ensure that policyholders will be able to receive their promised benefits up to a certain amount. The levels vary by each state but they all ensure a certain minimum amount. As of 2014, all state guaranty associations protect at least $100,000 in annuity benefits and $300,000 in life insurance death benefits (NOLHGA (2014)).

Since its creation, the NOLHGA has tracked the actual insolvencies that have occurred in the life insurance industry. In total, the NOLHGA has participated in 70 cases where an insurance company has been put into receivership. In spite of two recessions, including a financial crisis, the decade from 2000 – 2009 actually saw a reduction in the number of multi-state insolvencies compared to the 1990s. There were only 14 insolvencies from 2000 – 2009 but 25 from 1991 – 1994, with ten of those coming in 1994 alone. The insolvency time series for the multi-state failures is presented in Figure 2.3 Panel A (on next page), which shows both the raw number series and the insolvency ratio series. The insolvency ratio is simply the number of insolvencies in a given year divided by the number of life insurers in the industry in the prior year from ACLI (2014).

This variation seems to be due at least in part to the relative number of insurers in business at each time. The average number of stock and mutual life insurers was 1,883 from 1991 – 1994 and about 1,037 from 2000 – 2009, so about 1.3% – 1.4% of the total failed in each time period (ACLI (2014)). The multi-state insolvencies are also not evenly spread out geographically based on the state of domicile for the failed insurers and presented in Table 2.3. For reasons not explored in further detail here, Pennsylvania has experienced the most multi-state insolvencies with seven cases, but Alabama and Indiana are not far behind with five each.

Table 2.3 Life Insurance Company Failures by State

This table lists the states with the most cases of insolvency by life insurers both in terms of multi-state failures only (from 1991 through 2013) and all failures (from 1979 through 2013). Data are from the National Organization of Life and Health Insurance Guaranty Associations.

Multi-State Insolvencies		*All Insolvencies*	
State	**Number of Insolvencies**	**State**	**Number of Insolvencies**
PA	7	TX	65
AL	5	OK	20
IN	5	IL	19
MS	4	FL	18
TX	4	PA	18
CA	3	IN	14
IL	3	AL	12
OK	3	AZ	12
AZ, DE, FL, GA, ID, LA, MO, NC, NJ	2 each	CA	8
		NM	7

Figure 2.3 Time Series of Life Insurance Company Failures
These graphs track the year-by-year insolvency activity of life insurance companies. The multi-state insolvencies (those involving more than one state guaranty association) are in Panel A while Panel B contains the data for all insolvencies. For dataset, the graph includes a series for the raw number of insolvencies in each year ("Insolvency Number" on left vertical axis) from the National Organization of Life and Health Insurance Guaranty Associations and a series for the percentage of life insurers that fail in each year ("Insolvency Ratio" on right vertical axis). The percentage is the ratio of Insolvency Number and the number of companies in the life insurance industry at the end of the prior year from ACLI (2014). Note that the Insolvency Ratio for 1979 and 1980 in Panel B is based on the average number of life insurers from 1975 and 1980 as ACLI (2014) does not start providing annual figures until 1980.

Panel A – Multi-State Insurers Only

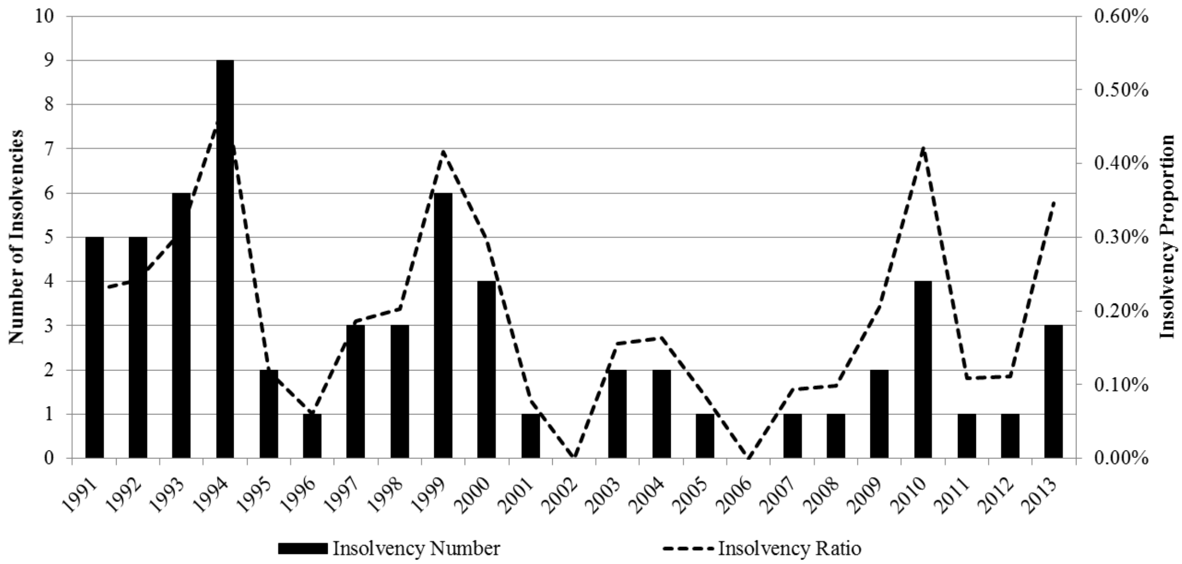

Panel B – All Insurers

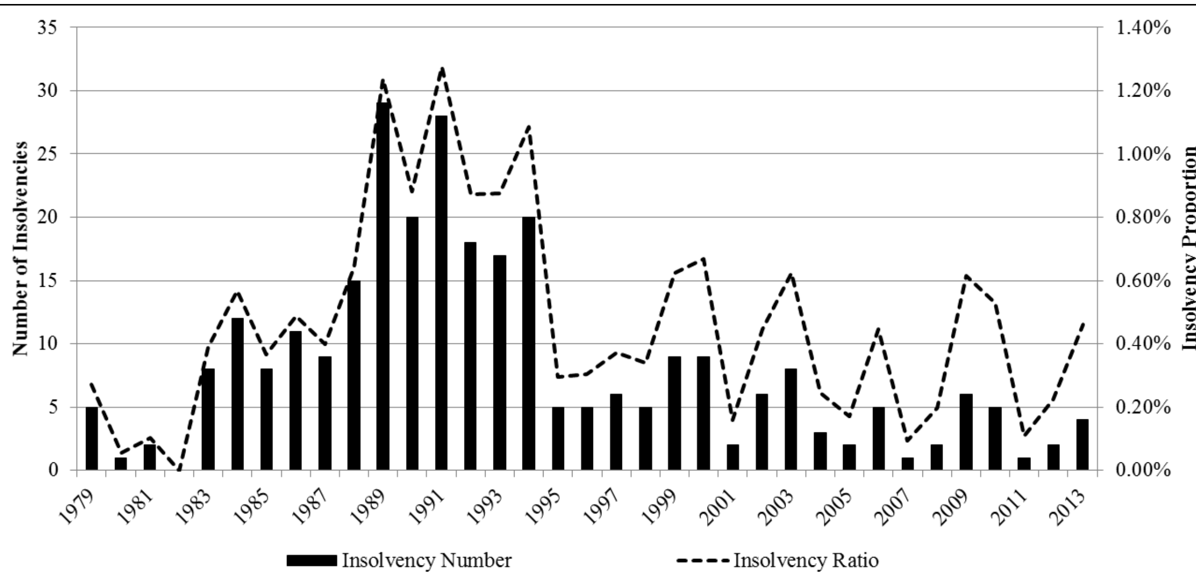

The NOLHGA also provides a more complete list of insolvencies that includes those involving only one state guaranty association and is presented graphically in Figure 2.3 Panel B. Surveying this list, which starts in 1979 compared with 1991 for the multi-state list, we derive similar conclusions. The 1991 – 1994 time period following the savings and loan crisis is again characterized by an elevated number of insolvencies with 83 cases out of 190 in total for all years. In contrast, there are only 20 insolvencies listed from 2008 – 2013, which is the period following an even greater financial crisis. In addition, we can now see that the early 1990s experienced a marked increase in insurer insolvencies as the insolvencies from 1979 through 1988, both in number and percentage terms, were much lower.

Once we include the smaller insolvencies, the state-by-state comparisons are somewhat different, though. Texas, which was tied for fourth in the multi-state list, is now the big winner (or loser) with 65 insolvencies occurring in that state alone followed by Oklahoma (20 insolvencies), Illinois (19), Florida (18), and Pennsylvania (18). Since 1979, there have been five states (Alaska, Maine, Nevada, Rhode Island, and Vermont) without any cases of insolvency. Interestingly, Connecticut has had only a single insolvency over this time period according to the NOLHGA in spite of being the traditional home of the "Insurance Capital of the World" (Hartford, from The Courant (2014)). Apparently, a large concentration of insurance business does not necessarily lead to a greater risk of insolvency, even in raw number terms.

F. Systemic Considerations[5]

The broader financial system has shown itself to be prone to occasional bouts of systemic problems that contagiously spread throughout the system. The 1800s experienced multiple financial crises, there was a financial crisis in 1907 that required the assistance of J.P. Morgan and his bank to stem the tide of banking failures in New York, the Great Depression involved bank runs that led to many bank closures and a subsequent reduction in available credit, and the recent financial crisis had ramifications across the global economy. There are several reasons that systemic risk in the financial system has the potential to inflict some wide-ranging and deep damage. Due to this destructive potential, financial regulators have started officially classifying certain financial institutions as being systemically important. The plan is to provide extra regulatory attention to these firms in order to minimize their systemic risk.

Traditionally, the financial system was very bank-based. Banks typically featured liquid short-term liabilities (deposits) and relatively illiquid long-term assets (loans). Although this maturity and liquidity transformation provides benefits to the economy, this also creates potential fragility. When a bank would experience a run on its deposits, it would need to call in loans, liquidate its loans at fire sale prices, or simply close. The contagion occurs because a run at one bank could incite runs at other banks as fears rise or as debtors of the first bank must withdraw their deposits elsewhere to meet the call-in of loans. Another source of risk is the fractional reserve nature of banking. By only holding a fraction of the deposits in reserve, the bank has created the potential of running through its reserves before it can meet all of the deposit demands.

Over time and especially in the past few decades, the financial system has become more market-based. Banks are still major financial institutions but others such as hedge funds, mutual

[5] The information in this section is credited to FRBNY (2006) unless otherwise cited.

funds, other asset managers, and brokers/dealers are now responsible for a greater proportion of the movement of funds throughout the economy. As a result, systemic risk has moved to the market as a whole rather than being focused on the health of specific institutions. In other words, systemic risk in a market-based system reveals itself more often through market-wide disruptions as opposed to being triggered by the demise of a specific entity. One benefit of the market-based system is that investment risk is more dispersed across many types of investors and institutions, rather than being concentrated in the savings and commercial banks.

However, it also has its own areas of weakness that enable systemic problems. The proper functioning of the financial markets depends on market marking and arbitrage activity. A market-based systemic event is often triggered by the significant decline, possibly unwarranted, in the price of some asset. The decline sustains itself when arbitrageurs, who would normally provide necessary balance to the market, are unable or unwilling to enter the market. As a result, market participants start selling this and other risky assets, which perpetuates the drop in asset market value. Hence, "market-based crises are often characterized by a coordination failure in which a wide cross section of participants in financial markets, including market makers and arbitrageurs, simultaneously decide to reduce risk taking and effectively pull back from financing activities." Although doing so helps each participant retain liquidity and capital, the aggregate effect on the market is sharp reduction in market activity and capital.

In either type of system, whether it is bank-based or the market-based, leverage is a factor that increases systemic risk. As financial institutions become more highly levered, the relative amount of capital available to cushion any unexpected losses dwindles. As a result, a relatively small decrease in the asset values could have a big impact on the capital base of a highly levered firm. If high leverage is widespread, it could lead to a broad sell-off of assets following a decline in prices and a systemic crisis. The widespread use of derivatives, even as a hedging strategy, could also increase the systemic risk of the financial system. The sensitivity of a derivative's value to underlying asset price can change as that underlying price changes. If selling the underlying asset is a reasonable response to the increased sensitivity, then systemic issues could result if many other market participants are trying to reduce their own risk in the same way.

III. Regulatory Framework for Life Insurance Companies

A. Purpose of Life Insurance Regulation

The purpose of life insurance regulation in the United States is stated clearly in a paper by the National Association of Insurance Commissioners ("NAIC").

> ***<u>US Insurance Regulatory Mission:</u>*** To protect the interests of the policyholder and those who rely on the insurance coverage provided to the policyholder first and foremost, while also facilitating an effective and efficient market place for insurance products (NAIC (2010)).

The NAIC is an association of the state-level insurance commissioners and has significant influence in the direction and shape of insurance regulation in the United States. Thus, the perspective of this important regulatory body is fixed squarely on the end user of insurance, which is the policyholder. In other words, the purpose of regulation, from the regulator's perspective, is not to maintain a steady and stable insurance market for its own sake but rather for the sake of protecting the persons who are depending upon the industry's products for their own financial security and stability. This mission provides the foundation for all of the regulations and policies proposed by the NAIC and other regulatory bodies.

B. United States Life Insurance Solvency Framework[6]

B. 1. The Uniqueness of the United States Solvency Framework

Life insurance regulation in the United States is somewhat unique because it has traditionally relied so heavily on state-level, rather than national or federal-level, regulation. This is largely due to the tradition of federalism in the United States. This was codified with the passage of the U.S. Constitution in 1787 and its subsequent amendments, particularly the tenth amendment. As a result, the federal government was given the power and authority to govern in specific areas largely related to handling issues between the various states and the country's foreign policy. Much of the remaining governing authority was delegated to the states. Thus, each state became responsible for regulating the insurance companies doing business within their particular jurisdiction.

Although "ultimate regulatory responsibility for insurer solvency rests with each state insurance department and the state insurance Commissioner," some degree of uniformity and consistency has developed over time through the activities of groups such as the NAIC. The NAIC offers "financial, actuarial, legal, computer, research, market conduct, and economic expertise to state regulators" and proposes model laws and regulations for state legislators and regulators to adopt. It has developed a solvency framework for insurance regulation in the United States that describes the regulatory system used in the United States and the core principles underlying it.

The NAIC describes ways in which the regulatory system in the United States is unique as a consequence of being state-based. It features "extensive systems of peer review, communication and collaborative effort that produce checks and balances in regulatory oversight" and a "diversity of perspectives with compromise that leads to centrist solutions."

[6] The information in this section is from NAIC (2010) unless otherwise cited.

Another feature of the solvency framework in the United States is that it is risk-focused. This means that regulatory “attention is paid to the greatest risks faced by insurers and the insurance market.”

Naturally, the first of these features likely is a result of the state-based system. With regulators residing in each state, as opposed to one primary regulatory body for the whole country, an opportunity is created for more collaboration and peer review to occur because there are many regulators operating across the whole country at any given time. This is also an important feature for a state-based system to remain effective as the life insurance industry has become more consolidated and as company operations cross more state lines. It does not require that all of the states craft the same regulations, but it is certainly helpful if the regulators of each state in which a particular company conducts business are able to collaborate when needed.

Regarding the second of these features, the NAIC believes that the U.S. system includes such a diversity of regulatory perspectives that, through compromise, it is able to avoid both of the regulatory extremes. These extremes are over-regulation, which “can impose unnecessary costs on consumers,” and under-regulation, which “can allow unnecessary harm to consumers and taxpayers.” Although the boundaries of each of these extremes is left vague, the NAIC undoubtedly believes that its model laws and regulations reflect a center-point between them since it is an association of the chief state regulatory officials who presumably fall along the full continuum of opinions on regulatory matters.

B. 2. Core Principles of the United States Solvency Framework

The NAIC lays out seven core principles of financial solvency that build on the regulatory mission to guide insurance laws and regulations. The first of these principles is regulatory reporting, disclosure, and transparency. This principle requires that "US regulators receive required financial reports from insurers on a regular basis that are the baseline for continual assessment of the insurer's risk and financial condition." This principle allows regulators to have information needed to measure the severity of a particular company's risk profile and its degree of financial distress. As such, it supports the regulator's monitoring role and gives it a signal for any potential regulatory action.

The second principle is off-site monitoring and analysis, which simply means that the "US regulators and the NAIC conduct off-site risk-focused analysis of insurers." There is where the regulators use the information provided by the first principle to conduct on-going analysis and monitoring of the key risks faced by life insurers. In addition to the information provided by the company, the regulator may use other publicly-available information, such as SEC filings, or information collected by the regulator in prior examinations.

The third principle is on-site risk-focused examinations where "[i]nsurers are subject to full-scope financial examination at least once every [five] years." In the second principle, regulators conduct on-going and high level monitoring of the insurer and its key risks. In the third principle, the regulator periodically conducts a detailed examination that analyzes many aspects of the insurer's business including corporate governance, management oversight, financial strength, risk identification and monitoring, and compliance with legal requirements. If the off-site monitoring indicates the need, then regulators may conduct these on-site examinations more frequently than five years or may do an on-site examination focused on a specific risk. At a minimum, though, regulators are required to do a full-scope on-site examination at least once every five years for each insurer in their jurisdiction.

The fourth principle is reserves, capital adequacy, and solvency where "insurers are required to maintain reserves and capital and surplus at all times and in such forms so as to provide an adequate margin of safety." As mentioned earlier, various types of reserves are the primary liability for a life insurance company, and these actuarially-determined reserves relate to the insurer's expected future payments to policyholders. Capital, or surplus, is essentially the difference between the value of an insurer's assets and its liabilities. Capital adequacy requirements seek to ensure that life insurers have a sufficient capital base to act as a cushion should the company experience unexpectedly high claims. In that case, the capital base "cushions" the hit to assets so that, hopefully, the company will not become insolvent.

In the United States, the primary set of capital adequacy requirements for life insurers is the risk-based capital ("RBC") system. It uses a standardized formula that "provides for higher RBC charges for riskier assets or for riskier lines of business so that more capital is needed as a result." There are certain thresholds that the risk-based capital base for each insurer is compared with to determine if the insurer is weakly capitalized and that regulatory action is required. One potential issue with the risk-based capital system is that it depends upon the charges to adequately measure the risk of each type of asset. For example, government debt is traditionally viewed as being very safe and so would receive a low risk-based charge. However, the European sovereign debt issues of the past several years cast into doubt the universality of that assumption across time and location. Unless the risk-based charges are updated in a timely and accurate fashion, it could actually result in insurers becoming more, not less, risky.

In addition to capital adequacy requirements, states also set minimum reserve requirements that ensure the insurer has the funds available to meet their obligations to policyholders under normal circumstances. The NAIC argues that setting minimum requirements on both reserves and capital bolsters the financial solvency of insurers in the United States by assuring that “policyholder obligations are covered by enough resources to meet most future economic scenarios and there are enough resources so that an adverse trend can be detected in time for the regulator to suggest/take corrective action.”

The fifth principle is regulatory control of significant, broad-based risk-related transactions and activities so that “certain…transactions/activities affecting policyholders’ interests must receive regulatory approval.” In other words, insurers may be barred from taking certain actions that fall outside the scope of routine underwriting and insurance issuance activities without first receiving explicit regulatory approval. The types of activities that merit such special scrutiny include licensing requirements, change of control, the amounts of dividends paid, transactions with affiliates, and reinsurance. Like other industries, major mergers and acquisitions may require regulatory approval that the transaction would not significantly impair the competitiveness of the industry or the welfare of policyholders. However, even dividends may be restricted although some states only require that extraordinary dividends require regulatory approval.

The sixth principle is preventive and corrective measures, including enforcement, where “the regulatory authority takes preventive and corrective measures that are timely, suitable and necessary to reduce the impact of risks identified during on-site and off-site regulatory monitoring.” The first four principles work to help the regulator identify and monitor the key risks that could potentially endanger the solvency of a particular insurer. This principle gives the regulator the authority to take actions that may be necessary to reduce the probability of insolvency. Of course, these actions are only necessary if the insurer itself has failed to properly mitigate them or take preventive measures on their own. The actions that could be taken include "requiring the insurer to provide an updated business plan in order to continue to transact business in the state; requiring the insurer to file interim financial reports; limiting or withdrawing the insurer from certain investments or investment practices; reducing, suspending or restricting the volume of business being accepted or renewed by the insurer; ordering an increase in the insurer's capital and surplus; ordering the insurer to correct corporate governance practice deficiencies; requiring a replacement of senior management; and seeking a court order to place the company under conservation, rehabilitation, or liquidation."

The seventh principle is exiting the market and receivership where the "framework defines a range of options for the orderly exit of insurers from the marketplace" if the insolvency occurs in spite of the best efforts of the company and/or the regulators. These options seek to minimize the damage of insurance company insolvency to policyholders and other creditors. These options include "mergers, acquisitions, reinsurance arrangements, non-renewal of part or all of the insurer's book of business, and…allowing the insurer to be placed in run-off mode under its own management."

C. Life Insurance Regulatory Framework Globally[7]

Just as the NAIC is a voluntary association of the insurance commissioners from the various states in this country, there is an international voluntary association of insurance regulators, the International Association of Insurance Supervisors ("IAIS"). It is an organization of insurance regulators from about 140 countries so it much if not nearly all of the global insurance market is influenced by this organization. Also similar to the NAIC, the IAIS's mission is "to develop and maintain fair, safe and stable insurance markets for the benefit and protection of policyholders and to contribute to global financial stability." So, it is clear that the regulators, at least officially, are on the side of the policyholders, which implies that regulatory risk is at least potentially one of the more important risks faced by life insurers. Given such a mission, regulators appear to be primarily concerned about the welfare of policyholders suggesting that the long-term prospects of any individual insurer are important only to the extent it furthers that primary concern.

The IAIS also propose core principles to support its regulatory guidance and activities, but in this case, it has 26 principles compared to seven for the NAIC. So, we will not go through them in much detail but will review them and highlight how they are similar and different than those of the NAIC. Principles one through three relate to the insurance supervisor in a particular regulatory jurisdiction. "The authority (or authorities) responsible for insurance supervision and the objectives of insurance supervision [must be] clearly defined" and the supervisor meets certain qualifications such as operating in an "independent, accountable and transparent" way and protects confidential information. The NAIC does expect these conditions as well, but it took them to be a precondition for effective regulation rather than core principles. The third principle sets an expectation that the regulators will exchange "information with other relevant

[7] The information in this section is credited to IAIS (2013) unless otherwise cited.

supervisors and authorities subject to confidentiality, purpose and use requirements." Principles 25 and 26 relate to this expectation that the local regulator will cooperate with other regulators and regulators in other jurisdictions when necessary. Recall that a unique feature the U.S. insurance regulatory system is the collaboration between the various state-level regulators, so this principle is naturally met without the need of the NAIC to explicitly state it.

The fourth principle simply states that insurers must be licensed prior to commencing insurance operations and that the licensing requirements must be "clear, objective, public, and be consistently applied." This principle is covered in this country with the NAIC's fifth principle which includes licensing requirements as one type of a significant, broad-based risk-related transaction or activity that requires regulatory approval. The IAIS sees this principle as an initial layer of policyholder protection, "a jurisdiction controls through licensing which entities are allowed to conduct insurance activities within its jurisdiction." In addition, the thirteenth principle of the IAIS, which discusses setting standards for reinsurance and other types of risk transfer, is covered by the NAIC's fifth principle as the NAIC includes reinsurance as another activity that is subject to regulatory review and approval.

Principles five through eight all relate to the upper-level management and governance of an insurance company. Although the NAIC's on-site examinations include reviewing the insurance company's corporate governance, it appears that the IAIS principles could be taking it a step further. The fifth principle requires that the regulator ensure that senior management, directors, other key persons, and even significant owners are and remain to be suitable. This principle could be either relatively innocuous or restricting depending on how the regulator interprets and defines suitability in each of these roles. Thus, it opens up an opportunity for bureaucratic decision-making in insurance regulation. The seventh principle also lays out in much more detail than the corresponding NAIC principle the necessary elements of an insurer's corporate governance framework, even proposing the necessary duties of board members and how boards should delegate some of its activities. One possibility for the greater detail in regulatory guidance is that corruption in corporate management is a bigger issue internationally than has traditionally been the case in the United States. Otherwise, principles six and eight simply require that transactions that involve significant changes in control receive regulatory approval, similar to the NAIC, and that insurers must have effective risk management systems and internal controls in place, which the NAIC would review with its periodic examinations.

Principles nine through 12 and 20 relate more directly to the actions of the regulator as it conducts insurance supervision. Many of these principles closely correspond to certain NAIC principles. The ninth and twentieth principles express the need for the regulator to take a risk-based approach, conduct both off-site monitoring and on-site examinations, and collect the information needed for these supervisory activities, which match up well with the first three core principles of the NAIC. In the event that supervisory activities uncover areas of concern, the tenth and eleventh principles charge the regulator with imposing and enforcing any necessary preventative and corrective measures, which corresponds to the NAIC's sixth principle. Ultimately, if actions by the insurance company and the regulator are insufficient to save the insurer, the twelfth principle, which corresponds to the NAIC's seventh principle, addresses the need for insurers to exit from the market and wind-down their insurance operations. This must be done in a way that "gives priority to the protection of policyholders and aims at minimising disruption to the timely provision of benefits to policyholders." Principles 23 and 24 elaborate further on these guidelines by stating that supervision should occur on a group-wide basis when the insurance company operates within a corporate group and that market-wide or economic environment factors should be utilized when monitoring and examining an individual insurer.

Principles 14 through 17 address issues of solvency and capital adequacy, on which NAIC's fourth principle is built. In this case, the NAIC adds more detail to the guidelines proposed by the IAIS. The IAIS proposes that regulators establish requirements for the valuation of assets and liabilities (principle 14), the investment activities of insurers (principle 15), enterprise risk management (principle 16), and capital adequacy (principle 17). Although the NAIC does not necessarily expand on these principles greatly in the statement of their own principles, the NAIC adds specific guidelines, such as risk-based capital calculations, to put the IAIS's relatively broad guidelines into practice. At a minimum, though, the IAIS does require that assets and liabilities be treated from an economic, rather than accounting, perspective for solvency and capital adequacy purposes. This means that the valuation of assets and liabilities must reflect the risk-adjusted present values of their cash flows and off-balance sheet investments may need to be included in the analysis.

The remaining principles relate to certain activities of the insurance company. Principle 18 requires regulators to also supervise insurance intermediaries, such as brokers who sell the insurer's products to customers, "to ensure that they conduct business in a professional and transparent manner." Principles 19 and 21 simply require that regulators ensure insurers treat policyholders fairly and effectively work to "deter, prevent, detect, report and remedy" fraud. Principle 22 requires that insurance companies put in place measures that will combat money laundering.

D. Banking Regulation vs. Life Insurance Regulation

In some ways, banking regulation and life insurance regulation are similar. Both have a large focus on the capital adequacy of the regulated institution, examine and monitor financial institutions on a regular basis, and take a risk-focused approach. For example, the Office of the Comptroller of the Currency ("OCC"), the primary regulator of banks operating with a federal charter, takes an integrated risk-based approach in which "the most significant risks to the bank will receive the most supervisory attention" (OCC (2008)). However, there are very important distinctions that must be made when discussing the regulation of insurers and banks.

Banks operate under a regulatory framework with a different structure than do insurance companies. While insurance regulation in the United States is state-based, banks operate under a dual banking system. It is a dual system in the sense that some banks operate with a state charter and are regulated by the corresponding state banking regulatory authority and other banks operate with a federal charter and are regulated by federal regulators. Thus, banking regulation is definitely more nationalized in the sense that federal agencies directly regulate a significant portion of the industry while another part of the industry is more directly regulated by the states.

Bank regulation also differs because the nature of banking is different than selling insurance products. A bank's traditional business model is to collect deposits from many diverse individuals and institutions who could demand to withdraw their full deposited amount at any time and then to make loans (many of which are long-term) using those deposits. Thus, it is inherently more susceptible to bank runs than insurers are to policyholder runs. Although policyholders may be building up an annuity account value or cash surrender value in an insurance policy, they may not be able to withdraw their funds at any time without paying a surrender charge. Even if they could, insurance policyholders are not as likely to run on an insurer as depositors. Bank runs occur because of a depositor's fear that others will withdraw their money from a failing institution first and then the money will run out. With life insurance products, though, the primary concern of the policyholder is that the insurer is capable of paying the scheduled annuity benefits and promised death benefits. Policyholders are not going to try to beat each other in line to face the grim reaper in order to get their money first when an insurer's finances are looking wobbly. Even if they did, insurers primarily invest in assets such as equities and investment-grade bonds, which are typically much more liquid than bank loans.

Another distinction between banking regulation and insurance regulation is related to the relative risk exposures of banks and life insurers. Banks can be very exposed to broad economic risk. When the economy goes into a recession, a bank's borrowers may start falling behind in their loan payments as a result of the poor economic conditions. This undermines the capital position of a bank. In contrast, poor economic conditions have no effect on a promised death benefit amount or on guaranteed annuity payments. In fact, if policyholders must cancel their insurance policies because they can no longer afford the premium payments, then the insurance company's capital position may actually be improved. They get to keep all of the premiums to date and no longer need to fund a future policy payout. That being said, some of the newest insurance products, those of the variable annuity and variable life insurance variety, do have exposure to the equity market and broader economic conditions. Although any promised payouts remain the same if the stock market falls, the amount that can be funded by the policyholder's account value likely does fall meaning the company itself must pick up the slack. Still, the shock mainly hits the insurer on a time value of money basis rather than creating an urgent liquidity crisis. The actual payments affected may be years in the future whereas a bank is losing cash flow now when economic conditions deteriorate.

E. Macro-Prudential Issues[8]

Macro-prudential regulation is becoming vastly more important now for financial regulation than in the past. The Dodd-Frank Wall Street Reform and Consumer Protection Act ("Dodd-Frank Act") of 2010 included a requirement that financial regulators adopt a macro-prudential approach. So, it is important to understand what macro-prudential regulation is and what some of the tools used by regulators are as this may have an impact on the regulation of and the management of risk by life insurers.

Financial regulation, both of banks and insurance companies, has traditionally focused on the health and stability of individual institutions or markets. Macro-prudential regulation instead focuses on systemic risk and minimizing the risk that financial disruptions drag down the broader economy. It builds this aggregate level perspective on top of the foundation achieved by traditional micro-prudential regulation. Acharya (2011) argues that traditional regulation may be insufficient to the extent that a particular institution's systemic risk is not internalized by that institution and the costs borne by other parties. In other words, systemic risk becomes a negative externality that is not factored in the decision-making of the entity causing the risk. Macro-prudential regulation focuses on the risks to the financial system in aggregate and whether certain risks are building up dangerously in the financial system.

There are several ways in which regulation has changed to identify systemic risks in the financial system. Institutionally, the Dodd-Frank Act has created the Financial Stability Oversight Council to identify risks to the financial system's stability, the Office of Financial Research to improve the quality of the information and data available to regulators for measuring and managing systemic risks, and other countries have created similar regulatory bodies to

[8] The information in this section is credited to Acharya (2011), Bernanke (2011), and Tarullo (2013) unless otherwise cited.

monitor systemic risks in their own financial systems. One of the powers of the new Financial Stability Oversight Council is the aforementioned practice of naming certain financial institutions as systemically important. These are firms that could supposedly cause significant turmoil in the financial system and/or broader economy should any of them fail. As a result, firms that receive this designation become subject to extra regulatory attention and scrutiny.

There are a range of potential macro-prudential tools available to regulators. Tarullo (2013) gives some broad classifications of macro-prudential tools. Some may be termed "lean-against-the-wind" measures because their intent is to prevent systemic risks from building up in the financial system while others are called resiliency measures because they seek to make firms and the system more resilient if systemic risk accumulates and manifests itself anyway. Another classification given by Tarullo (2013) is time-varying and time-invariant where the distinction depends on whether a measure is "turned-on" based on an increase in systemic risk or is always on regardless of the risk levels seen throughout the system.

It appears that federal financial regulators plan on focusing their macro-prudential regulatory efforts on the resiliency and time-invariant class of tools. Daniel Tarullo, a member of the Board of Governors of the Federal Reserve System, has stated that "building greater resiliency is central to the macro-prudential regulation of large financial institutions" and time-varying measures will have a more limited role. He also mentions some examples of the resiliency measures used already or at the disposal of the Federal Reserve and other regulators. Stress testing the largest banks, initially conducted in 2009, was one of the first macro-prudential tools used by the Federal Reserve after the financial crisis. This is one example of a resiliency measure that can be both time-varying and time-invariant. A stress test is time-invariant to the extent that it is done on a regular basis regardless of the systemic risk conditions in the financial system. However, the hypothetical scenarios used or the risk weights of various asset classes could be time-varying to reflect experience and new conditions from one year to the next. For example, the Federal Reserve modified the market shock scenario in the 2011 round of stress tests to account for the European sovereign debt crisis. Also, knowing that the risk weights would reflect their relative exposure in some hypothetical scenarios can create some "lean-against-the-wind" disincentives that encourage firms to avoid loading up too much in assets particularly exposed to stress conditions.

Another resiliency measure being used by the regulators is identifying certain bank and non-bank financial institutions as being systemically important and then subjecting them to extra scrutiny. These systemically important institutions are those that could inflict great damage to the financial system or the broader economy should they fail. The extra capital requirements for these systemically important firms will be macro-prudential by building in extra buffers for the negative externalities that would result if any of the other systemically important firms failed. Extending this designation to non-bank financial institutions is apparently one way of increasing the coverage of the macro-prudential regulation. However, there is not a full consensus on how this is being implemented, even within the regulatory community. For example, Prudential Financial and MetLife, two major American life insurers, were designated as systemically important financial institutions in September 2013 and December 2014, respectively. In both cases, the insurance-specific members of the council opposed the designation essentially arguing that the other council members failed to understand the distinctions between banking and insurance (FSOC (2013) and FSOC (2014)). They argued that insurers are not subject to the same run risk as banks for similar reasons as those outlined in Section III.D. Nonetheless, the council decided to go ahead and subject an insurance company to direct federal regulation, which itself is unprecedented in the United States. In response, MetLife has sought judicial review by a federal court to have the designation rescinded (Business Wire (2015)).

Time-varying macro-prudential measures may have limited efficacy due to some practical difficulties. As of now, there is no consensus on a reliable systemic risk measure or set of measures. Acharya (2011) outlines some of the possibilities, but the literature has not coalesced around one that could be used for regulatory purposes. Even if there was a clear method of measuring systemic risk, there may be timing issues. For example, Tarullo (2013) cites how the Basel III framework includes a countercyclical capital buffer of up to 2.5% that could be "turned on" when systemic risk is building but also gives banks up to a year to meet the revised capital requirement. Given the one-year waiting period, the build-up in systemic risk could result in a crisis situation before the bank has changed behavior or acquired extra capital.

IV. Risk Management Tools and Strategies for Life Insurance Companies

A. Traditional Asset Allocation

The oldest method used to manage the inherent risks of the life insurance business is through traditional asset allocation. By allocating their assets appropriately, life insurers seek to maximize their returns while minimizing the probability that their reserves and capital surplus will be insufficient to fund the benefit payments. The typical asset classes that have been used in this risk management strategy are stocks, bonds (especially government bonds and long-term corporate bonds), mortgages, and real estate. Although policy loans used to comprise a more significant percentage of total assets (nearly 12% in 1920 according to ACLI (2014)), this would not be considered an investable asset class in this context because the amount that gets "invested" in policy loans is driven by policyholder borrowing decisions rather than the investment policies of the insurer.

As we have mentioned earlier, life insurers use traditional asset allocation to manage their risks by investing primarily in long-term investment-grade bonds. This investment provides better duration and cash flow matching with the long-term insurance guarantees than other types of investments. Investing in mortgages and real estate can also support this long-term investment policy. Insurers have traditionally avoided giving significant allocations to the equity market as this asset class features significantly more market price volatility than fixed income investments.

Basically, life insurers are looking to sell long-term financial stability and bonds provide more asset stability than stocks do, so life insurers invest in bonds and avoid stocks. The equity market is not avoided entirely as it has higher expected returns than bonds, but the allocation has been kept to a limited level to minimize the equity market risk. An interesting challenge for life insurers in recent decades as investment-type policies (such as variable annuities and variable life insurance) have grown in popularity has been managing the risk of their Separate Account assets. If the insurers seek to minimize their exposure to equity market risk, policyholders seem to flock to it. We saw earlier how 80% or more of the Separate Account assets have been invested in stocks in recent decades. Because these policies may come with certain guarantees, the life insurer ultimately takes on at least some of this equity market risk even though they have limited control over the initial investment decision. We will discuss in more detail how insurers can manage this additional exposure to equity market risk in the next section.

In addition to focusing their investments in certain asset classes with favorable risk characteristics, life insurers are selective in how they allocate their monies within an asset class to better manage their risk. For example, life insurers do not put all of their bond investment into the lowest-rated tier of investment-grade bonds even though that would increase their expected returns. Instead, they seek to optimize their risk-return profile by investing across a mix of rating tiers.

Table 4.1 analyzes the credit profile of the life insurance industry's bond and mortgage investments for the same years shown in the asset distribution of Table 2.2. The NAIC classifies the credit risk of bonds into several classes with bonds of the highest quality going into Class 1 and those of the lowest quality going into Class 6 (this allows them to include both publicly-traded and privately placed bonds). Life insurers almost exclusively invest in bonds of the highest qualities with only 5.75% of investments in 2013 going to bonds in medium or low quality classes. Comparing this with 1999 and 2008 suggests that this behavior has been consistent over recent history, although the percentage in bonds of less than high quality has been trending down. As with stocks, life insurers focus most of their investing activity in lower risk investments but still look for some opportunity to enhance expected returns by not completely avoiding the riskier choices.

Table 4.1 Credit Profile of Life Insurer Bond and Mortgage Holdings

This table classifies the aggregate bond and mortgage holdings of life insurance companies into various credit risk categories. The bond profile is in Panel A while the mortgage profile is in Panel B. For bonds, the risk categories correspond to the various classes defined by the NAIC. Classes 1 and 2 are deemed to be "High Quality," Class 3 is of "Medium Quality," and Classes 4, 5, and 6 are of "Low Quality." For mortgages, the categories are defined according to the current status of the mortgage. The data for 1999 and 2008 are from ACLI (2010), and the data for 2013 are from ACLI (2014).

Panel A - Bonds			
	1999	*2008*	*2013*
High Quality			
Class 1	64.58%	67.68%	62.61%
Class 2	28.08%	26.01%	31.62%
Medium Quality			
Class 3	4.15%	3.62%	3.63%
Low Quality			
Class 4	2.67%	1.72%	1.56%
Class 5	0.43%	0.78%	0.42%
Class 6	0.10%	0.19%	0.14%

Panel B – Mortgages			
	1999	*2008*	*2013*
In Good Standing	97.89%	99.84%	99.46%
Restructured	1.80%	0.06%	0.46%
Overdue	0.18%	0.04%	0.05%
Foreclosed	0.14%	0.06%	0.04%

Life insurers appear to be even more concerned about credit quality in their mortgage portfolios. Since the height of the internet boom of the late 1990s, nearly all of their mortgages have consistently been paying in full and on time. Even during the recent financial crisis (when concerns about the credit quality of mortgages underlying mortgage-backed securities contributed to market-wide stresses) and several years later, a very small percentage (0.10% in 2008 and 0.09% in 2013) of their mortgage investments are either in "overdue" or "foreclosed" status. This appears to be largely due to the risk characteristics of the mortgages chosen for investment. For example, nearly 96% of the mortgage investments have loan-to-value ratios of 80% or less and over 86% of all mortgages have loan-to-value ratios below 71%. Hence, life insurers apparently seek to manage their credit risk by focusing almost all of their mortgage investments in mortgages of very low credit risk.

B. Hedging Programs

As more of a life insurer's asset allocation decision-making is transferred to policyholders (through the Separate Account assets), the need arises for the company to use methods other than traditional asset allocation to manage risk. Because the insurer only has indirect control over the asset allocation decisions of policyholders, it cannot simply adjust its weights to various asset classes to reduce its exposure to certain risks such as equity market risk. In theory, the company could eliminate stock investment options from its Separate Account policies, but based on the revealed preferences of policyholders, that would likely be very unpopular. One method of managing the company's exposure to increased risk as a result of these policyholder allocation decisions is to implement a hedging program.

Hedging programs make use of derivative securities to achieve a desirable risk profile without placing too many constraints on the policyholder investment options. Life insurers would naturally be exposed to a sudden down shock in the equity market given the high proportion of Separate Account assets invested in stocks and the presence of certain guarantees on those policies. If the equity market returns are high, or even just mediocre, then the policyholder's account value can fund most or all of the guaranteed benefits. However, if equity market declines, then the account values may no longer be sufficient to fund benefits payments and creating a liability for the company in expected present value terms. To offset or hedge this risk exposure to equity market declines, a life insurer could short equity futures contracts or purchase equity put options, so that the hedge program produces profits at the same time policy liabilities increase. Of course, hedging programs are not cost-free. When no such shocks to the equity market occur, the put option premiums paid or the losses on futures contracts produce a real cash outflow from the company.

As with many financial decisions, there is a trade-off between hedging away the additional equity market risk and "self-insuring," or building up extra reserves within the company to protect against any adverse equity market movements. Of course, life insurers need not choose one option to the exclusion of the other. They can certainly use a mixed strategy that protects against adverse equity shocks up to a certain level such as 20% and then hedges away the exposure to worse shocks using derivatives. One benefit of this is that it cheapens the hedging program by using derivatives that are significantly out-of-the-money when the trades are executed.

C. Reinsurance[9]

Traditional asset allocation and hedging programs focus on the asset side of the insurance company's balance sheet. They take the insurance liability risks as a given and seek to manage them by investing in a mix of assets that is expected to generate cash inflows that match up well with the expected liability cash outflows and maximize profit without undermining the insurer's long-term stability. In contrast, reinsurance seeks to manage the liability risks by transferring at least some of the risk to another party. Although reinsurance may show up as an asset on an accounting balance sheet, the life insurer is effectively reducing its potential liability by purchasing its own insurance policy. This risk management strategy is also very common. As of 2013, 89% of life insurers receiving life insurance premiums purchased at least some reinsurance to better manage their risks.

Reinsurance is insurance for an insurance company whereby risk is transferred from a cedant (the insurance company) to a reinsurer. Reinsurers tend to be multi-national firms so global diversification of life insurance risks occurs through reinsurance. Life insurers essentially guarantee a certain portion of their policyholders' tail risk (where the tail refers to either the left (traditional life insurance) or the right (life annuities) tail of the policyholders' remaining life probability distribution), and reinsurers then guarantee a certain portion of the life insurer's tail risk. Thus, the reinsurer is the ultimate guarantor of the policyholder's benefits corresponding to the portion of the life insurer's business covered by reinsurance. As a result, managing risk through reinsurance may increase the life insurer's counterparty risk but with the benefit of adding greater certainty to the expected cash outflows. Another benefit is that reinsurance may reduce the life insurer's regulatory reserve or capital surplus requirement.

[9] The information in this section is credited to Wehrhahn (2008) and ACLI (2014) unless otherwise cited.

Reinsurance serves multiple potential purposes for the life insurer seeking coverage. By transferring some of the risk to the reinsurer, the life insurer can increase the amount of insurance coverage provided than would otherwise be possible based on their own financial strength. Or, it allows the insurer to cover a particularly unique and/or large risk that would not be prudent to take on in isolation. An expanded capacity to underwrite insurance also allows the insurer to diversify the risks acquired by using their additional capacity to enter new markets, lines of business, or regions.

Another use of reinsurance is akin to the role that the life insurer itself plays for the policyholders. By seeking insurance coverage, policyholders are often trying to protect themselves and their families from a low-probability but catastrophic event. Similarly, the life insurer can manage the risk of a catastrophic event by transferring it to a reinsurer. In general, this is probably less of a concern to the life insurance industry since it is very rare for even natural disasters to result in mass loss of life. Nonetheless, an insurer with operations focused in a place such as Florida may seek some protection from the possibility that an event such as a hurricane leads to an unexpectedly high number of claims. This also shows how reinsurance can serve the purpose of allocating work to those with the competitive advantage to handle it. Reinsurers end up covering the risk of various potential catastrophes across many different areas of the world. As a result, they tend to build up a high level of expertise in pricing, forecasting, and underwriting the risk of such events. By transferring it to the reinsurer, the life insurer is effectively utilizing the reinsurer's relative expertise in covering this risk.

Insurance companies can even use reinsurance as a source of financing for a new line of business. A reinsurance agreement could be made where the reinsurer provides the insurance company with future expected profits of the new business as a reinsurance commission. The insurance company must then pay the reinsurer out of the actual profits of the new business over time. Since reinsurers have more insurance-specific expertise and understand the nature of the risks better, this form of financing may be cheaper and more effective than using other forms of financing such as a bank loan.

Reinsurance agreements between the cedant and the reinsurer can also take several different forms. The general classification is between proportional and non-proportional reinsurance. Proportional is the case where the reinsurer and the insurance company each cover some specified percentages of the risk (and corresponding premiums, expenses, reserves, etc.) under consideration. This includes a quota share arrangement where the same percentage is transferred for each risk included in the agreement or surplus/excess of retention reinsurance where each risk is shared in different proportions. Even if the covered risks have different sharing proportions, each of them would include a certain risk level (called the retention or surplus line) that the insurance company is willing to cover on its own and then the reinsurer would cover the proportion above this threshold up to a maximum amount (called the capacity). Risk amounts in excess of the capacity are not insured in such an arrangement.

Non-proportional reinsurance is used when insurance companies are looking to protect themselves from adverse impacts resulting from a spike in actual claims (as opposed to the ex-ante amounts at risk). One example of this type of reinsurance is excess of loss or stop loss reinsurance, which is similar to the surplus proportional agreement except now we are dealing with actual losses sustained by the insurance company. The cedant will cover losses up to some threshold (called the priority in this case), the reinsurer will cover losses up to the capacity, and losses above the capacity are uncovered by the agreement. These thresholds can be determined on the basis of actual dollar amounts or percentages of premiums (e.g., the priority is 80% of the total premiums, the capacity is 40% of the total premiums, and anything above 120% of the total premiums is not insured). These agreements could potentially include a feature where the cedant must pay a copayment for losses in excess of the priority.

The remaining types of non-proportional reinsurance are excess of time and catastrophe agreements. An excess of time agreement is related to the excess of loss type but is more conducive to certain insurance policies (such as disability or long-term care insurance) that require the insurer to pay out recurring payments for potentially an extended period of time. The reinsurance arrangement allows for the life insurer to manage the risk of making payments much longer than expected by having the reinsurance cover some or all of the payments after some specified period of time. Catastrophe agreements are similar to the excess of loss type in that the actual losses must exceed the priority for it to be covered by reinsurance. The distinguishing feature of catastrophe agreements is that the losses are due to some specified catastrophic event as opposed to the insurer experiencing an unexpected amount of potentially unrelated claims over a time period.

D. Catastrophe Bonds

If an insurance company is particularly exposed to catastrophe risk, which is the risk that some catastrophic event could conceivably produce a spike in claims, then issuing catastrophe bonds can be a method of managing the financial implications of that risk. They do this by freeing up funds for benefit payments that would otherwise be paid out to creditors.

The advantage for risk management provided by catastrophe bonds comes from the conditional nature of its cash flows. At issuance, lenders purchase bonds from the insurance company as with any other bond. The key distinction is that the bond payments to creditors are contractually dependent on whether or not a specified catastrophe has occurred. If the catastrophe has occurred, then the company is no longer obligated to re-pay the creditors. If no catastrophe occurs, then the bond payments continue until maturity just like any other standard corporate bond. In this way, the insurer has managed its exposure to this risk by transferring it to the broader capital markets. Obviously, the catastrophic event that triggers nullification of the remaining bond payments is precisely specified in the bond contract. Although it introduces a type of basis risk for the insurer where a catastrophe could occur that generates a lot of claims but does not qualify for the bond trigger, it also enables a deeper market for the catastrophe bond (creditors would be less willing to take on a vague catastrophe risk when natural disasters occur somewhere on a regular basis).

Although this risk management tool is used by life insurers much less than the tools mentioned thus far, it could conceivably be useful to a life insurer under certain situations. If a life insurer is operating in a region that is particularly exposed to a recurring risk of potentially deadly disasters (such as Florida for hurricanes, certain areas of the Midwest for tornados, California for earthquakes, and certain Asia-Pacific areas for tsunamis), then catastrophe bonds could help mitigate even a low probability of financial distress should catastrophe occur and result in a widespread loss of life. They could even help protect the company against the relatively new risk of terrorist attacks in North America and Europe.

E. Government Regulations

Although ultimately outside the control and certainly not often welcomed by the life insurer, government regulations are a method by which the insurance company's risks could be managed or mitigated. In fact, regulations of financial firms often have risk mitigation as their primary goal. Although we already engaged in a more detailed discussion of insurance regulation, some key points are relevant here as well.

For life insurers, the major regulations related to risk are the reserves and capital regulations governing the necessary funds that firms need to have on hand to maintain long-term solvency. Ensuring long-term financial stability flows from the regulatory mission covered earlier of protecting the policyholders since insolvency can cause major disruptions for policyholders even if the state guaranty association ultimately makes them whole. Recall that the state guarantees cover insurance benefits only up to a certain level, though. As a result, some policyholders may lose benefits if the insurance company fails even if they have paid all of the required premiums. To minimize these liquidity and insolvency risks, states and nations require that insurers meet certain regulatory thresholds on reserves and capital surplus. Since the state has decided to guaranty at least some of policyholders' benefits, the state also has a vested interest in reducing the cost of such a guaranty program by increasing the solvency of life insurers.

Regulations also seek to manage other risks than the risk that insurance companies are failing to keep sufficient funds on reserve. They also try to mitigate risks arising from how insurers might invest those funds. Insurers do not invest primarily in investment-grade bonds entirely on their account. State insurance regulations often restrict the amount of weight that can be given to higher risk asset classes such as equities and high-yield ("junk") bonds. These investment constraints serve to manage the insurance company's market and credit risk by not allowing it to load up their investment allocations in financial securities highly exposed to these risks.

To be sure, government regulations are not an infallible risk management tool. At times, they can actually generate greater risk of financial distress. By relying on the wisdom of supposedly prudential regulations, private enterprises (including life insurers) can develop an attitude of taking advantage of any opportunity to boost profits within the boundaries of the regulations. For example, the risk charges of various assets for risk-based capital calculations may not always keep pace with the inherent riskiness of the corresponding assets. Insurers will naturally prefer to invest in a bond with higher yield but the same risk charge as another. As a result, insurers may end up exposing themselves to assets (e.g., mortgage-backed securities in 2007 or European sovereign debt in 2008-2009) that appeared to be safe according to bond ratings and regulatory risk charges but were ultimately revealed to be much riskier. Thus, a potential weakness of using government regulations as a risk management tool is the possibility that they will encourage firms to sacrifice their own due diligence in order to maximize potential profits (i.e., moral hazard).

F. Interrelationships among the Various Tools

Perhaps unsurprisingly, an insurance company's use of one or more of these risk management tools can impact the need or effectiveness of another tool. This is due, at least in part, to an overlap in the risks being addressed by the various risk management strategies. Asset allocation, hedging programs, and government regulations are all focused on managing the market risk of the company's investments. Credit risk is also managed primarily by asset allocation and government regulations. Insurance risks are addressed by reinsurance agreements and catastrophe bonds. Systemic risks are increasingly being managed through government regulations, particularly for any life insurers who become designated as systemically important financial institutions.

One example of the linkages between the various tools is the following relationship between asset allocation and hedging programs. Life insurers have increased their exposure to the equity market largely by granting certain policyholders the ability to make their own investments using the policy's account value. As we have described earlier, the historical experience thus far has been that policyholders with these Separate Account assets tend to give large asset allocation weights to stock-related investments. The company could manage this increased equity exposure by re-adjusting its own allocations to traditional asset classes within the General Account. It could increase even higher its allocation to safer and more stable bond investments, or it could avoid purchasing any speculative debt as these bonds may be exposed some of the same residual claim risks as equities. However, by implementing a hedging program, the company can minimize the need to alter its own General Account asset allocations to balance the increased equity exposure from the Separate Account. It can largely offset this exposure with a hedging program that generates payoffs from derivative securities at the same time as equity market risks increase Separate Account-related liabilities. Although a hedging program is not cost-free (e.g., the premiums required to purchase put options), it can help the company avoid needing to reduce its General Account market risk (and thereby reducing the expected returns) as a result of greater Separate Account market risk.

We already alluded to another potential linkage between government regulations and asset allocation as risk management tools. As governments take some responsibility for maintaining the long-term safety and soundness of life insurers, a natural tendency is for the companies themselves to start relaxing their own prudential asset allocation practices in order to maximize expected returns and profits while still meeting the regulatory requirements. Government regulations can also reduce the need to use the other tools. As we have discussed earlier, a significant component of the life insurance regulatory framework focuses on the capitalization of life insurers. As regulations require insurers to maintain larger capital buffers, the probability falls of an insolvency-inducing spike in claims or market value movement. As a result, insurance companies may have less incentive to purchase reinsurance, issue catastrophe bonds, and implement hedging programs. Thus, the relationships between these tools are not always self-reinforcing.

Ultimately, an interrelationship between all of these risk management strategies that they are trying to manage the ultimate risk of insurance company failure, and a life insurer fails when it cannot pay the promised benefits to policyholders at the scheduled times. Thus, all of these tools seek to minimize liquidity risk while balancing the need to provide required rates of return to owners. Asset allocation and hedging programs do this by protecting the company's assets from debilitating losses in market values. Reinsurance and catastrophe bonds do this by providing for additional financing or risk sharing when realized claims exceed the resources provided by the company's insurance pricing and forecasting models and start to undermine the capital buffers. Government regulations do this by ensuring that the company is sufficiently capitalized to weather some unexpected losses, has sufficient funds in reserve for scheduled payouts, and is managing their assets appropriately.

V. Measuring and Modeling Life Insurance Market Risk

A. Statistical Measures of Risk

The notion that many human activities are exposed to uncertainty and risk is certainly not new. The earlier discussion about the early history of life insurance makes this clear. The existence of ancient burial societies and conditional loans to sailors would not make sense unless sailors and soldiers were aware that undesirable events might occur to them. However, the concept of quantifying and mathematically measuring our risk exposure is not so old.

Although we are not going to argue that no statistical measures of risk were being used prior to Markowitz (1952), that paper has certainly had a large impact on our measurement of risk, especially in the context of forming investment portfolios. The measure of risk used in that paper was the variance, or standard deviation, of returns. Given a sample of returns for any asset, the variance is calculated as a weighted average of squared deviations from the mean. Unfortunately, the variance is not easily interpreted given that it is measured in percent-squared. The standard deviation, which is simply the square root of the variance, is much easier to interpret and has often been used to measure portfolio risk since Markowitz (1952).

The implicit assumption in using variance and standard deviation, though, is that risk is any volatility in potential outcomes relative to the expected outcome. In other words, returns both 1% above and 1% below the expected return contribute equally to an asset's variance. Naturally, investor attitudes about both of these possibilities are not so symmetrical. Investors are much more concerned about the possibility that returns will be below the expectation and may not even consider volatility on the other side to be risk at all. So, a number of risk measures have been developed that focus on measuring "downside risk," or the likelihood that one will be left disappointed by an especially poor outcome.

One of these downside risk measures is called the second lower partial moment and has a similar calculation as the standard variance except for one key thing. Only returns that are below some target rate set by the practitioner contribute to the lower partial moment. The remaining returns do not count for the risk calculation. We can then calculate the square root of the lower partial moment if we desire to have a risk measure with the same units as our returns.

By considering only returns below some threshold, we are focusing our risk measure on the left tail of the probability distribution of returns. Other risk measures that focus on the left tail include conditional tail expectations, which estimate the expected outcome conditional on being in the left tail (i.e., when things go bad, how bad exactly do we expect it to get) and value-at-risk, which estimates the minimum amount of loss should we find ourselves in the left tail. One can also fully model the distribution of returns conditional on being in the tail by fitting a probability distribution the returns in the left tail. Doing so may provide one with a fuller picture of an asset's downside risk by trying to capture more of the return behavior in the left tail than a single risk measure.

B. Importance of Tail Risk

Focusing on tail risk is especially important in the context of this study in managing market risk for a life insurance company. As we discussed earlier, life insurers must focus on their own long-term strength and stability given the long-term nature of their liabilities. Thus, life insurers need to be keenly aware of their exposure to tail risk. Such risk can be much more ruinous to the solvency of the company than the normal day-to-day market volatility. If not managed well, the manifestation of tail risk can quickly wipe out much or all of a company's surplus capital, at which point the regulatory authorities will likely require corrective action up to and including the termination of the company as a going concern. Given that even a single manifestation of tail risk has the potential to produce these consequences, life insurance companies tend to invest conservatively and sacrifice some expected return by focusing on bonds rather than stocks in their General Account investments.

Traditional diversification strategies may be insufficient for the purpose of managing tail risk, though. This is because the financial markets have a tendency to "behave as one" in the words of Junior and De Paula Franca (2012). They and others have observed that correlations of asset returns across a number of asset classes are appreciably higher during crisis periods than during other periods. Unfortunately, these markets are correlating in the exactly the wrong direction from the investor's perspective. In other words, when a crisis hits, many of the markets go down together. As a result, increasing one's allocation to a normally safe asset class such as high-grade corporate debt may not be enough to avoid trouble when tail risk rears its ugly head. An exception may be investing in U.S. Treasury securities to the extent they are perceived to be a safe haven during distress. For example, U.S. Treasury bonds, as tracked by the Bank of America/Merrill Lynch U.S. Treasury indices, provided returns of over 10% from September

2008 through March 2009 whereas corporate bonds lost nearly 8%, commercial mortgage-backed securities lost over 18%, and stock investors saw nearly half of the value of their holdings wiped out.

This shows that it is both vital and difficult for a long-term investor such as a life insurance company to measure and manage its exposure to tail risk. Not doing so puts at risk the financial security of numerous policyholders who are also the company's owners if it is a mutual insurer. As a result, this study will analyze the asset allocation decision from a "downside risk" perspective and will base this decision on returns adjusted for downside risk. This will be done by making use of the Sortino Ratio, which was developed by Brian Rom of Investment Technologies according to Booth and Broussard (2015). This ratio is similar to the commonly used Sharpe Ratio, with a couple of key exceptions. The numerator is the difference between the return of the asset under study and a target or minimum acceptable rate determined by the researcher or practitioner. The denominator is a measure of downside risk, which measures the volatility of the returns below the minimum acceptable rate.

The measure of downside risk used in this study is the second lower partial moment. As mentioned earlier, this risk measure is calculated similarly to variance except only observations below some threshold or target rate of return count towards the measurement. Thus, it is based on a philosophy that risk is the potential for significant losses to occur rather than volatility on either side of an expectation. This fits the perspective of an investor like a life insurance company that may sustain significant hits to capital from adverse financial market movements but will not be affected so much by the opposite.

C. Probability Distributions to Model Tail Risk

C. 1. Generalized Pareto Distribution

One of the probability distributions often considered to model the tail returns of an asset (see Castillo and Hadi (1997), Longin (2005), and Booth and Broussard (2015)) is the Generalized Pareto Distribution ("GPD"). Using this distribution implies one makes a certain assumption about the nature of the tail. Namely, it assumes that the tail returns are all of the returns below some threshold set by the researcher regardless of which time period they come from. As a result, if one's dataset includes a time period with particularly poor returns, such as the September 2008 – March 2009 time period for many risky assets, then many of those returns will be included in the tail even if they were not relatively worse than other returns from the same time period. This lends the GPD to be useful in many downside risk frameworks where one's risk measure is based on returns below some minimum acceptable rate.

Fitting the GPD to data requires the estimation of three parameters. One of which, the location parameter (λ) or the threshold, is set by the researcher. The other parameters model the scale (σ) and tail shape (τ) of the distribution. The threshold and shape parameters have $(-\infty, \infty)$ as their support while the scale parameter is non-negative. In the context of this study, though, the threshold will be chosen in order to focus on the left tail of asset returns.

Depending on the value of τ, the GPD may simplify into either the exponential or the continuous uniform distribution. If $\tau = 0$, then it is equivalent to an exponential distribution with a mean equal to σ, and if $\tau = 1$, then it is equivalent to a uniform distribution on the range $[0, \sigma]$. In addition, $\tau > 0$ corresponds to a fat-tailed distribution and $\tau < 0$ corresponds to a finite distribution without a tail (see Longin (2005)). Otherwise, the cumulative distribution function $F(R)$ of the GPD given return R is

$$F(R) = 1 - \left[1 + \tau\left(\frac{R-\lambda}{\sigma}\right)\right]^{-\frac{1}{\tau}}$$

and the probability distribution function *f*(*R*) is

$$f(R) = \frac{1}{\sigma}\left[1 + \tau\left(\frac{R-\lambda}{\sigma}\right)\right]^{-\left(\frac{1}{\tau}+1\right)}$$

Because the distribution is defined is terms of observations above a threshold, $R - \lambda$, the returns used in this analysis are all multiplied by -1 so that the left tail will be the portion of the sample above the threshold.

The GPD theory, however, does not definitely give an estimate for the location of the tail (λ). Following Booth and Broussard (2015) who base their threshold selection method on Loretan and Phillips (1994), we will base our estimate of the threshold on the empirical rule that $k = n^{2/3}/\ln[\ln(n)]$ where k is the number of observations included in the tail and n is the sample size. We will then estimate the GPD under four cases, which include having k / 2, k, 2k, and 4k observations in the tail, and select the threshold producing the best fit.

Using the GPD to model tail risk enables one to calculate the lower partial moment with a probabilistic and ex ante approach (see Booth and Broussard (2015)). This makes our estimate of downside risk more conducive to a portfolio allocation problem where the focus is on what may happen rather than explaining what did happen from an ex post perspective. Rather than calculating it as

$$LPM = \frac{1}{T}\sum_{i=1}^{T}\left(R_{target} - R_i\right)^2, \forall\, R_i \leq R_{target}$$

where T denotes the number of observations in the left tail as defined by R_{target}, we can make use of the theoretical distribution of R. Therefore, we will calculate the lower partial moment to be

$$LPM = \int_{-\infty}^{R_{target}} \left(R_{target} - R\right)^2 d[f(R)]$$

where $f(R)$ corresponds to the probability distribution function of the estimated GPD model.

C. 2. Generalized Extreme Value Distribution

According to Castillo and Hadi (1997), the traditional alternative to the GPD approach is to model tail observations with a Generalized Extreme Value Distribution ("GEVD"). The GEVD also includes an implicit assumption about the nature of the tail, but it is a different one than that of the GPD. The GEVD assumes that the tail is defined in relativistic terms by including only the most extreme observations from each block of observations or time period. In other words, the tail is the set of block maxima (or minima) where the size of the block is set by the researcher. This ensures that the tail used for the analysis equally represents each time period of the sample. However, it also means that the tail could include observations that would not appear to be very extreme in the context of the full sample.

Like the GPD, the GEVD is based on the location, shape, and scale parameters (see Singh (1998)). The cumulative distribution function $F(R)$ is given as

$$F(R) = \exp\left(-\left[1-\tau\left(\frac{R-\lambda}{\sigma}\right)\right]^{\frac{1}{\tau}}\right)$$

and the probability distribution function $f(R)$ is given as

$$f(R) = \frac{1}{\sigma}\left[1-\tau\left(\frac{R-\lambda}{\sigma}\right)\right]^{\frac{1-\tau}{\tau}} \exp\left(-\left[1-\tau\left(\frac{R-\lambda}{\sigma}\right)\right]^{\frac{1}{\tau}}\right)$$

As with the GPD, returns that are modeled with the GEVD are multiplied by -1 in order to fit within the block maxima approach. We can also estimate the lower partial moment with the GEVD by replacing the $f(R)$ in the probabilistic calculation with the probability distribution function of the estimated GEVD model.

D. Models to Manage Tail Risk

D. 1. Risk Hyperplane

Ultimately, in order to say something meaningful about the tail risk borne by a life insurer as a result of their asset allocation decisions, we need a method by which to analyze the risk of the company's whole portfolio of assets. To do this requires us to model the joint distribution of the life insurance company's assets. Under the Gaussian assumption for each of the individual assets, the process of modeling the joint distribution is relatively straightforward. However, it has been observed by many (e.g., Mandelbrot (1963), Fama (1965), Cont (2001), Longin (2005)) that this assumption fails to be robustly supported by the actual returns of these assets. Given the inherent difficulties of directly estimating a joint probability distribution when the underlying marginal distributions are not Gaussian, we will use alternative methods of modeling the joint behavior.

One of these methods will be a risk hyperplane approach. Given a set of asset allocation weights and time series returns for individual assets typically owned by life insurance companies, one can calculate a time series of joint portfolio returns and model the representative company's tail risk conditional on the initial weights. Instead of choosing between the two definitions of a tail, both the GPD and GEVD probability models are used to model the portfolio tail risk in this approach. By repeating this process for a whole range of initial weights, one can build a hyperplane based on the expected portfolio returns and tail risk across a set of feasible asset allocation choices for a life insurance company. The hyperplane created by this approach can then guide us in determining which one of these portfolio weight sets is optimal for the purpose of maximizing portfolio return while also placing a high premium on managing tail risk. The Sortino Ratio will be used to measure this trade-off and determine the optimal portfolio.

In order to define the feasible set of asset allocation choices for a representative life insurance company, we will make use of both the historical industry-wide asset allocation weights found in the ACLI Fact Books for various years and the regulatory constraints within which life insurers must make their decision. As a starting point, the portfolio weights will be based on the industry weights for the most recent year available, which is 2013. These weights vary from year to year and the full range of historical industry weights will be used to determine the set of feasible asset allocation choices. Some of the regulatory constraints that are particularly relevant for this exercise include restrictions on the proportion of General Account assets invested in equities and other “high-risk” securities.

D. 2. Copulas

Copulas will be a second method utilized to estimate the joint probability distribution of a life insurer's investment portfolio. Copula theory goes back to Sklar (1959), which lays forth the following theorem. Suppose one has random variables R_1, …, R_d that have continuous cumulative distribution functions F_1, …, F_d and a joint cumulative distribution function F. Sklar argues that there exists a unique copula C such that $F(r_1, \ldots, r_d) = C(F(r_1), \ldots, F(r_d))$. In other words, the copula function transforms the marginal cumulative distribution functions into the joint function F. Note that C is a distribution function on $[0, 1]^d$ with uniform marginal distributions, which allows one to use the distribution functions of R_1, …, R_d as the functional inputs. As a result, copulas hold out the promise of being able to estimate the joint probability distribution of returns from several asset classes even when returns are not Gaussian.

In practice, though, most of the copula functions that have been identified involve pairs of random variables. This poses a challenge for the joint analysis of three or more random variables. Work by Joe (1996), Bedford and Cooke (2001, 2002), and Kurowicka and Cooke (2006) makes clear a handy method of skirting this limitation as outlined in Brechmann and Schepsmeier (2013). These authors propose using vine copulas, which decompose a multivariate copula into a series of conditional bivariate pairs. This allows one to make use of the available bivariate copula functions for modeling a multivariate problem.

In order to make use of copulas for this study, we base our approach on similar copula-based portfolio modeling work done by others including Deng, Ma, and Yang (2011), Brechmann and Czado (2013), Allen, McAleer, and Singh (2014), and Carmona (2014). We first need to prepare our raw data to be useful as inputs into a copula function. To do this, we will estimate a GPD model for each of the life insurer asset classes in order to model their marginal distributions. In this case, we will use a two-tailed GPD model because we are using the vine copula to model the joint dependence between the asset classes across the full distribution. We will focus on the left tail at a later stage when we are dealing with the returns of a prospective portfolio. We have chosen to use the GPD model in our vine copula analysis, rather than the GEVD, since this is more consistent with the prior vine copula literature. The copula data are formed by taking the cumulative distribution function, $F(r_i)$, for each return r in asset class i based on i's GPD-based marginal distribution.

Next, an appropriate conditioning path is chosen based on the dependence structure of these assets, and an appropriate bivariate copula function is chosen for each "twist" of the vine. The structure selection criteria proposed in Czado (2010) and Czado, Schepsmeier, and Min (2012) are utilized for this purpose. The structure selection depends not only on the joint dependence of various bivariate pairs of the individual asset classes but also on the underlying structure of the vine copula. In particular, vine copulas can take on multiple structures that determine how the variables are paired together. One structure, called a canonical vine copula ("C-vine copula"), is one where each level of the vine (called a tree) has a single root variable and all pairs are built on this root (e.g., if we have random variables R_1, R_2, R_3, and R_4 and 1 is the root, then the pairs for the first tree would be 1-2, 1-3, and 1-4). Naturally, this is chosen when there is a variable in each tree of vine that drives a lot of the joint dependence. Another structure, called a drawable vine copula ("D-vine copula"), is one where no single variable drives the joint dependence throughout the vine or in each tree (e.g., if we have random variables R_1, R_2, R_3, and R_4 and 1, 2, 3, and 4 is selected to be the correct order for the first tree, then the pairs for the first tree would be 1-2, 2-3, and 3-4). For this dissertation, we will model the joint dependence of a representative life insurer's assets using both a C-vine copula and a D-vine copula structure and compare the results.

In the C-vine copula, the root variable for each tree is the variable that has the greatest joint dependence across the other variables. To do this, we estimate Kendall's tau for each pair of variables in our dataset and keep the absolute value of the estimate. Then, we sum up the absolute tau estimates for each variables, and the one with the greatest sum is chosen to be the root variable for this tree. After the first tree, the Kendall's tau estimates and bivariate copula selection is conditional on any root variables from prior trees, which are otherwise excluded from the analysis in this level of the vine.

In the D-vine copula, we are not looking for a root variable that drives the joint dependence in each tree but are instead looking for the variable order that will maximize the joint dependence of the first tree. The later trees are naturally built based on the order of the first tree (e.g., if we have random variables R_1, R_2, R_3, and R_4 and the pairs making up the first tree are 1-2, 2-3, and 3-4, then the pairs for the second tree are 1-3|2 (from 1-2 and 2-3) and 2-4|3 (from 2-3 and 3-4)). Again, we use Kendall's tau to measure the joint dependence for the purpose of selecting the correct variable order for the first tree. The variable order that produces the highest sum of tau estimates (using the absolute value of each estimate) for the first tree determines the correct variable order for estimating the vine copula.

To estimate the canonical and D-vine copulas, we use the CDVine package for the statistical software R. This package contains many functions useful for statistical inference, estimation, and analysis of canonical and D-vine copulas. Given the vine copula structure selected by the researcher, the software package selects the best choice of bivariate copula function for each pair in the vine and estimates the parameters of the copula function. It does this by estimating the parameters of each possible copula function choice for each pair and using the Akaike information criterion to select the best one. Prior to selecting the copula function, it will also test that the two variables are statistically independent, given the appropriate conditioning set if necessary. If the null hypothesis of independence cannot be rejected at the significance level chosen by the researcher, then the independence copula is chosen, which assumes no joint dependence at all at this point in the vine. For this dissertation, we use a significance level of 5% to test for independence.

After estimating a vine copula, we simulate returns for a typical life insurer's investable assets that reflect the joint dependence modeled in the vine copula and the GPD marginal distributions. To build portfolios of these returns, we use the same weights as in the risk hyperplane analysis. The difference here is that we are modeling the joint dependence theoretically using the vine copula rather than estimating the portfolio characteristics directly from the historical data. To compare the simulated portfolios, we estimate the Conditional Value-at-Risk, or expected shortfall, based on various threshold points as well as the Sortino Ratio from the risk hyperplane analysis. The expected shortfall is the first lower partial moment of the portfolio's asset return distribution while our measure of tail risk used to estimate the Sortino Ratio is the second lower partial moment. Ultimately, the Sortino Ratio will guide our selection of the optimal asset allocation for the General Account of our representative life insurance company.

E. Data

E. 1. Data Sources and Preparation

The historical asset allocation weights for the life insurance industry come from the ACLI Fact Books, which are published on an annual basis. We have access to Fact Books published in 2005 and 2008-2014 which gives us weights for 1994, 1997-2004, and 2006-2013. We are able to get data for years prior to 2004 because certain tables within the Fact Books categorize assets not only for the most recent year but also for the prior year and ten years prior for comparison purposes.

The 2014 Fact Book classifies the assets held by life insurers into the following asset classes: U.S. government bonds, non-U.S. government bonds, corporate bonds, mortgage-backed bonds, common stocks, preferred stocks, farm mortgages, residential mortgages, commercial mortgages, real estate, policy loans, short-term investments, cash and cash equivalents, derivatives, other invested assets, and non-invested assets. Those last two categories include assets such as premiums or investment income due to the company but not yet received. All of these categories were reviewed in light of needing to define a readily available proxy investment with sufficient historical data to merit inclusion in the study. As a result, not all of these categories will be included but a very significant portion of life insurer assets will be covered. In particular, U.S. government bonds, non-U.S. government bonds, corporate bonds, mortgage-backed bonds, common stocks, residential mortgages, commercial mortgages, short-term investments, and cash and cash equivalents will be included in the study. As of 2013, this accounts for nearly 84.6% of the General Account, 96.5% of the Separate Account, and 89.2% of the combined assets. The categories with the largest allocations that are not included are policy loans, other invested assets, and non-invested assets.

Table 5.1. Proxies and Data Availability for Includible Life Insurer Assets

Asset	*Proxy*	*Data Availability Date*	*Abbreviation*
U.S. Treasury Bonds	The BofA Merrill Lynch U.S. Treasury Composite Index	10/31/1986	*trbd*
Non-U.S. Treasury Bonds	The BofA Merrill Lynch Global Government Excluding the U.S. Composite Index	9/30/1993	*fnbd*
Corporate Bonds	The BofA Merrill Lynch U.S. Corporate Composite Index	10/31/1986	*corp*
Mortgage-Backed Bonds	The BofA Merrill Lynch US Mortgage Backed Securities Index	1/6/1989	*rmbs*
Common Stocks	CRSP Value-Weighted Index (with distributions)	10/31/1986	*vwst*
Residential Mortgages	The BofA Merrill Lynch US Mortgage Backed Securities Index	1/6/1989	*rmbs*
Commercial Mortgages	The BofA Merrill Lynch US Fixed Rate CMBS Index	12/31/1997	*cmbs*
Short-Term Investments	The BofA Merrill Lynch US 6-Month Treasury Bill Index	3/31/1992	*trbd6*
Cash and Cash Equivalents	The BofA Merrill Lynch US 3-Month Treasury Bill Index	3/31/1992	*trbd3*

Proxies were determined for each of the asset classes to be included in the study, and they are listed in Table 5.1. Daily index values for the Bank of American/Merrill Lynch indices are from Bloomberg. The daily value-weighted stock index total returns are from the Center for Research in Security Prices ("CRSP"). Daily index values were determined based on the total returns provided by CRSP. Although stock return data are available prior to October 31, 1986, this date corresponds to when data becomes available for asset classes that traditionally receive much greater allocations. As a result, only stock returns following this date are used. Also note that the composite indices for U.S. Treasury bonds, non-U.S. Treasury bonds, and corporate bonds are weighted averages of the individual Bank of America Merrill Lynch indices covering the following maturity ranges: one to five years, five to ten years, ten to fifteen years, and fifteen or more years. The weights for these averages are based on the maturity distribution for life insurers as of 2013 according to the 2014 ACLI Fact Book. This Fact Book provides industry allocations for the following maturity ranges: one to five years, five to ten years, ten to twenty years, and twenty or more years. Thus, the weights for the one to five year and five to ten year indices were set equal to those provided in the 2014 Fact Book. The weight for the ten to fifteen year index was assumed to be equal to half of the ten to twenty year allocation in the 2014 Fact Book. The weight for the fifteen or more year index was the other half of the ten to twenty year allocation plus the allocation of the twenty or more year range.

For all of the proxies, daily log prices were calculated based on the index values, and then log returns were calculated by taking first differences of the log prices. Descriptive statistics of these proxies are given in Table 5.2. Panel A includes the statistics for the full sample time period available for each variable, and Panel B includes the same statistics but on a common date range available to all variables (i.e., December 31, 1997 to December 31, 2014). The mean daily return, minimum daily return, and maximum daily return statistics are in percentage terms as a result of being multiplied by 100 (i.e., a mean of 0.025 for *corp* means that corporate bonds return 0.025% each day on average). Note that *cmbs* only has 4,234 observations over the common date range (rather than 4,236 observations as with the other variables) due to missing observations for two days during this time period.

Generally, a normal distribution would fail to precisely model these asset classes. Other than a few exceptions (*trbd* and *trbd6* in the full date range and *trbd6* in the common date range), they exhibit skewness significantly different from zero. Interestingly, they do not all exhibit skewness in the same direction. We find negative skewness for *trbd*, *corp*, *vwst*, and *cmbs* and positive skewness for *fnbd*, *rmbs*, *trbd3*, and *trbd6*. All of them are "heavy-tailed" since they have excess kurtosis that is significantly different from zero. Based on the Lagrange Multiplier test, these assets exhibit GARCH effects as all but *fnbd* in the common date range (p-value of 0.0117) are significantly different from zero at least at the $p = 0.0007$ level.

Table 5.2. Descriptive Statistics of Daily Data

Column headings correspond to an individual asset class. The point estimates for skewness and excess kurtosis are augmented with t-statistics based on the null hypothesis of zero skewness and excess kurtosis. The Ljung-Box and Lagrange Multiplier statistics are based on the null hypothesis of no linear dependence and GARCH effects, respectively. The statistics contained in Panel A correspond to the full data available for each individual asset class where as those in Panel B correspond to the date range that is common to all asset classes (12/31/1997 – 12/31/2014).

Panel A – Full Date Ranges								
	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
Mean (× 100)	0.027	0.025	0.029	0.038	0.026	0.024	0.012	0.013
Variance (× 100)	0.0022	0.0032	0.0010	0.0129	0.0004	0.0022	0.0000	0.0000
Skewness	-0.03	0.17	-0.29	-1.00	-0.13	-3.14	0.74	0.07
	(-1.17)	(4.93)	(-9.79)	(-34.27)	(-4.11)	(-83.35)	(22.92)	(2.10)
Excess Kurtosis	3.35	3.25	3.32	18.98	5.42	80.56	26.18	43.14
	(57.30)	(48.29)	(56.90)	(325.54)	(88.84)	(1070.07)	(403.20)	(664.31)
Minimum (× 100)	-2.737	-3.211	-2.492	-18.796	-1.850	-9.358	-0.255	-0.414
Maximum (× 100)	4.575	5.142	2.583	10.876	1.757	4.812	0.192	0.284
Ljung-Box	17.07	7.56	12.38	17.91	98.27	200.37	2973.79	914.36
Lagrange Multiplier	52.11	11.41	125.93	122.75	65.19	495.81	66.38	81.86
Observations	7,039	5,311	7,041	7,059	6,459	4,234	5,691	5,691
Panel B – Common Date Range								
	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
Mean (× 100)	0.024	0.023	0.025	0.026	0.021	0.024	0.009	0.010
Variance (× 100)	0.0025	0.0035	0.0012	0.0163	0.0004	0.0022	0.0000	0.0000
Skewness	-0.12	0.21	-0.34	-0.28	0.11	-3.14	0.89	0.04
	(-3.22)	(5.49)	(-9.07)	(-7.46)	(2.80)	(-83.35)	(23.73)	(1.07)
Excess Kurtosis	1.83	3.18	2.21	7.00	5.42	80.56	34.32	56.87
	(24.32)	(42.27)	(29.39)	(93.01)	(72.01)	(1070.07)	(455.97)	(755.50)
Minimum (× 100)	-2.737	-3.211	-2.492	-9.405	-1.362	-9.358	-0.255	-0.414
Maximum (× 100)	3.145	5.142	2.117	10.876	1.757	4.812	0.192	0.284
Ljung-Box	10.73	6.14	4.38	26.60	61.71	200.37	2144.47	726.84
Lagrange Multiplier	80.18	6.36	108.29	179.78	90.96	495.81	54.42	62.80
Observations	4,236	4,236	4,236	4,236	4,236	4,234	4,236	4,236

Given that we are concerned with the tail risk of a portfolio of assets in this study, we will also describe the data in terms of joint dependence. To do these we will use the concept of an exceedance correlation. As described by Patton (2004), an exceedance correlation of asset returns measures the correlation between two sets of returns conditional on both returns exceeding a certain percentile. For the downside case, this is expressed mathematically as

$$corr[X, Y | X \leq Q_X(q), Y \leq Q_Y(q)]$$

where X and Y are the returns for two assets, $Q_X(q)$ is the qth percentile of asset X, $Q_Y(q)$ is the qth percentile of asset Y, and q is less than or equal to 0.5.

For the sake of brevity, we will show the exceedance correlations with only the equity returns, but the concept could be applied to any pair of asset classes. We will also show how the joint dependence varies across time in our data sample by calculating the exceedance correlations for each year. To do so, we find each day within the year where both daily returns exceed their respective thresholds. In order to ensure there are sufficient observations within each year to calculate the correlation, we will use a fairly wide tail based on $q = 0.3$, or the thirtieth percentile.

These exceedance correlations are plotted in Figure 5.1. To provide a benchmark, we also plot an average of the full correlations ("Overall" on Figure 5.1) between equities and each fixed income asset class. To calculate this average, we first measure the correlation between equities and each fixed income asset class using all daily returns in the data sample. We then average these correlations in order to derive an average level of overall dependence between equities and the fixed income asset classes. In our data sample from January 1998 through October 2014, this overall correlation is about -0.14. From reviewing Figure 5.1, we can see that correlation between equities and various types of fixed income securities tends to increase as you move to the left tail of returns. This is not always the case, of course, as certain years and asset classes, especially corporate bonds in 2011, actually exhibit less joint dependence in the left tail than overall. The time period from 2000 to 2003 also has generally less joint dependence between equities and fixed income in the left tail than in the rest of the data sample. Still, it is clear that joint dependence with equities is generally higher during periods of tail risk events. Although this generally coincides with the observations of Hong, Tu, and Zhou (2007) and Junior and De Paula Franca (2012) cited earlier, it is interesting to note that these exceedance correlations do not seem to be particularly high during years of known market stress such as 2002 or 2008.

Figure 5.1. Exceedance Correlations with Equity Returns

This graph plots the exceedance correlations of the fixed income asset classes with equity returns. They are calculated for each year shown on the horizontal axis. An exceedance observation is a day on which the asset's return is more negative than that same asset's thirtieth percentile for that year. Given each fixed income class, the exceedance correlation is calculated as the correlation between the returns from all days where both equities and the fixed income asset have an exceedance. The "Overall" measurement is an average of the full correlations between equities and each fixed income asset class.

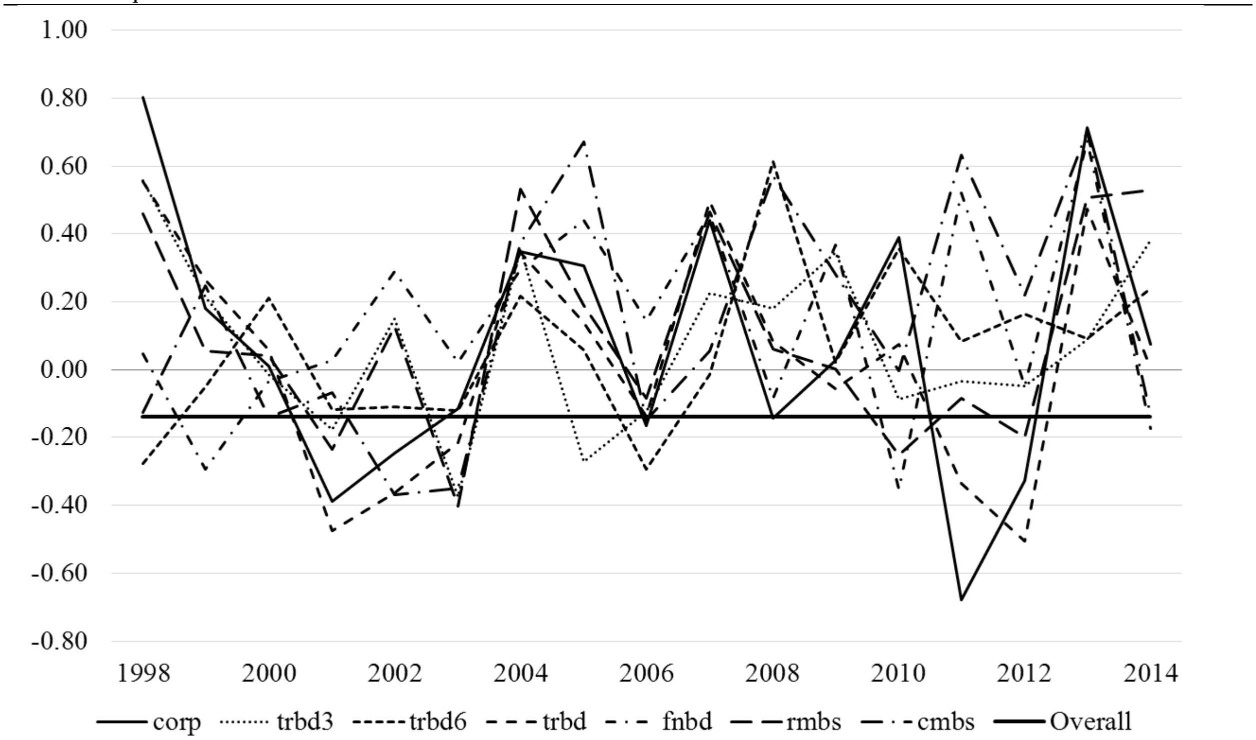

E. 2. Unit Root and KPSS Tests

Before we estimate the marginal distributions of the includible life insurer assets, we test for the presence of unit roots in the prices and returns of these assets. Log prices and returns are used for these tests to correspond to the data that will be used to estimate the marginal distributions. A unit root is a time series concept and relates to the autocorrelation of one observation with the observation from the prior time period. In general, the relationship between two time series observations of successive time periods can be expressed as

$$y_t = \rho y_{t-1} + u_t$$

where y_t is the observation at time t, ρ is the one-lag autocorrelation in this time series, and u_t is the error at time t. Testing for a unit root in this time series is equivalent to testing the null hypothesis that $\rho = 1$.

We transform the basic setup as Dickey and Fuller (1979) by subtracting y_{t-1} from both sides to get

$$\Delta y_t = \gamma y_{t-1} + \varepsilon_t$$

where y_t is the log price or log return at time t and $\gamma = \rho - 1$. Now, testing for a unit root is equivalent to testing the null hypothesis that $\gamma = 0$. Failing to reject this null hypothesis suggests the presence of a unit root in y while rejecting the null hypothesis means that no unit root is present.

Our results of this unit root test are presented in Table 5.3. According to Hamilton (1994), the critical values to reject the null hypothesis at the 10% and 1% levels of significance are -5.7 and -13.8. Reviewing the *t*-statistics for the regression in prices, we certainly fail to reject the null hypothesis of a unit root in log prices for every asset class. When testing returns, though, the point estimates become significantly negative with *t*-statistics much greater in magnitude than even the critical value for the 1% level of significance. These results suggest the presence of unit roots in log prices but not in log returns. Given that log returns are the first differences of log prices, this indicates that the prices are I(1) variables and the returns become I(0) after differencing.

Table 5.3. Unit Root Test Results

For each asset class listed in the column headings, daily changes in log prices or log returns are regressed against the prior day's observation. The coefficient on the prior day's observation ($\hat{\gamma}$) is given below for each asset class and type of data (prices or returns). The *t*-statistics are in parentheses.

y	*trbd*	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
$\hat{\gamma}_{prices}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(4.84)	(3.15)	(7.52)	(2.66)	(9.76)	(3.34)	(55.03)	(43.83)
$\hat{\gamma}_{returns}$	-0.985	-0.981	-0.953	-1.006	-0.915	-0.847	-0.456	-0.591
	(-82.60)	(-71.48)	(-79.97)	(-84.52)	(-73.83)	(-55.74)	(-40.99)	(-48.88)

In addition to testing for unit roots in our data, we also conduct a Kwiatkowski, Phillips, Schmidt, and Shin (1992) ("KPSS") test on the log asset returns for stationarity. This test uses a null hypothesis that a particular time series is stationary around a deterministic trend rather than a null of a unit root as in the Dickey-Fuller tests. It uses a Lagrange Multiplier statistic to test this null hypothesis.

The test is conducted by regressing y_t, $t = 1, 2, \ldots, T$, on an intercept term and a time trend. The residuals e_t from this regression are saved and used to calculate two numbers. The first, S_t, is the sum of the residuals from time 1 through time *t*. The second, $\hat{\sigma}_\varepsilon^2$, is an estimate of the error variance for *y*, and it is equal to the sum of squared residuals divided by *T*. The LM test statistic is equal to

$$LM = \sum_{t=1}^{T} S_t^2 / \hat{\sigma}_\varepsilon^2$$

Asymptotically, this becomes

$$\hat{\eta} = T^{-2} \sum S_t^2 / s^2(l)$$

where $s^2(l)$ is a consistent estimator of the long-run variance σ^2.

Table 5.4. KPSS Stationarity Test Results

For each asset class listed in the column headings, a Kwiatkowski, Phillips, Schmidt, and Shin (1992) ("KPSS") test is conducted. This test relies on a Lagrange Multiplier statistic ($\hat{\eta}$) calculated with the residuals from a regression of y on an intercept and time trend. Point estimates are provided below, and the *p*-values are in parentheses.

y	*trbd*	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
$\hat{\eta}$	0.018	0.043	0.031	0.060	0.067	0.059	0.905	0.775
	(0.9886)	(0.6780)	(0.8582)	(0.4523)	(0.3782)	(0.4641)	(<0.0001)	(<0.001)

The results of the KPSS stationarity tests are provided in Table 5.4. With the exception of the short-term U.S. Treasury Bills, we clearly fail to reject the null hypothesis of stationarity around a time trend for each of these asset classes. Again, the Dickey-Fuller tests for *trbd3* and *trbd6* strongly suggested the lack of a unit root in these daily log returns, so we are confident that all of these time series variables are suitable for the subsequent analysis.

F. Estimation of Marginal Distributions

F. 1. Full Date Range Available

To better understand the behavior of left tail returns for the set of includible life insurer assets, we model the marginal distribution of each asset class with both a GPD model and a GEVD model. We first fit these models to the asset classes over the entire time series available for each asset and then repeat the estimation for the pared-down date range that is common to all assets. The "evir" package in R was used to fit the GPD and GEVD models to these time series variables.

Fitting these models requires the researcher to set a particular threshold (for the GPD model) or block size (for the GEVD model). This is not entirely objective as the relevant definition of the tail can vary by application and researcher. Using the Loretan and Phillips (1994) method for the GPD, though, provides with a more systematic way of selecting an appropriate threshold.

For the GEVD estimations, we estimate the model four times using various block sizes related to different time periods. In particular, we estimate the model using a block size of five (a week), 21 (a month), 126 (half of a year), or 252 (a full year). That being said, some of the variables were not estimable at lower block sizes, so the set of block sizes was adjusted for these cases on a case-by-case basis. The variables for which the estimations were adjusted in this way include *trbd3* (used ten observations instead of five), *trbd6* (used ten observations instead of five), *rmbs* (used 75, 126, and 189 observations instead of five, 21, and 126), and *cmbs* (used 21 and 63 observations instead of five and 21).

Table 5.5. Marginal Distribution Estimation over Full Date Range

The table contains the marginal point estimates and observation counts for the includible life insurer asset classes. The GPD estimates are in Panel A, and the GEVD estimates are in Panel B. The *t*-statistics are included in parentheses. The number of extreme observations or the block size chosen for each asset class is also included. The GPD and GEVD models were estimated several times for each asset class with a varying number of tail observations or block sizes. The particular set of tail observations used is based on the number of total observations in the series (from Loretan and Phillips (1994)) and thus depends on the individual asset class. The block sizes were chosen from a set of five, 21, 126, and 252 observations. Some of the variables, such as *rmbs*, used alternative block sizes if one or more of the standard set produced erroneous estimates.

Panel A – Generalized Pareto Distribution Estimates

	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
$\hat{\lambda}$	-0.0054	-0.0061	-0.0035	-0.0119	-0.0021	-0.0031	-0.0000	-0.0001
$\hat{\tau}$	0.0020	-0.0124	0.0545	0.2034	0.1080	0.4415	0.5030	0.3356
	(0.05)	(-0.31)	(1.69)	(4.76)	(3.33)	(8.61)	(4.89)	(5.43)
$\hat{\sigma}$	0.0031	0.0036	0.0021	0.0074	0.0014	0.0019	0.0001	0.0001
	(27.96)	(22.84)	(68.17)	(18.38)	(694.28)	(946.68)	(32.79)	(55.02)
Tail Observations	672	568	672	672	640	496	148	296
Total Observations	7,039	5,311	7,041	7,059	6,459	4,234	5,691	5,691

Panel B – Generalized Extreme Value Distribution Estimates

	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
$\hat{\lambda}$	-0.0066	-0.0042	-0.0043	-0.0128	-0.0040	-0.0036	-0.0001	-0.0002
	(-38.64)	(-36.42)	(-37.59)	(-29.08)	(-20.33)	(-22.03)	(-27.47)	(-106.37)
$\hat{\tau}$	0.0496	0.0015	0.0967	0.2494	0.1293	0.3028	0.6619	0.4661
	(1.22)	(0.07)	(2.17)	(5.47)	(1.42)	(4.14)	(5.42)	(3.18)
$\hat{\sigma}$	0.0030	0.0034	0.0021	0.0072	0.0018	0.0026	0.0001	0.0002
	(72.35)	(78.34)	(1049.46)	(23.23)	(906.83)	(1280.92)	(38.39)	(88.10)
Block Size	21	5	21	21	75	21	126	126
Total Observations	7,039	5,311	7,041	7,059	6,459	4,234	5,691	5,691

The GPD and GEVD models chosen for each variable are presented in Table 5.5. Panel A contains the location (λ), shape (τ), and scale (σ) parameter estimates along with *t*-statistics for the relevant estimates for the GPD models while Panel B contains this information for the GEVD models. Reviewing these results, we see that there is a range of threshold and scale estimates as a result of variability in the daily volatilities of these asset classes. For example, *vwst* and *fnbd* have tails that are located relatively farther from zero while those of U.S. Treasury bills are very close to zero. The tail shape parameters for some of the asset classes, especially *trbd*, *fnbd*, *corp,* and *rmbs,* are either not significantly different from zero or only marginally so.

This suggests that the tail risk of these variables may be modeled appropriately by an exponential distribution with mean $\hat{\sigma}$. In contrast, the scale parameters are always very significantly different from zero in a statistical sense and generally in the range of 0.0015 to 0.0035. However, *vwst*, *trbd3*, and *trbd6* are notable exceptions with *vwst* having a much larger scale estimate and the U.S. Treasury bills having much lower estimates.

F. 2. Common Date Range

Due to differences in data availability across the set of includible asset classes, we also estimate the marginal distributions on a pared-down date range that is common to all assets, which is January 2, 1998, through October 31, 2014. This common date range includes 4,236 trading days for all assets, although *cmbs* is missing observations on two occasions.

When estimating the GPD models, the number of observations to include in the left tail is based on the model chosen for that variable on the full date range. For example, Table 5.5 shows that the left tail for *corp* includes 672 observations over the full date range out of 7,041 observations in the full time series. Thus, the left tail over the common date range was chosen to include the $4{,}236 \times (672 / 7{,}041) \approx 404$ most extreme observations from this date range.

When estimating the GEVD models, the block size was chosen to be the same as the model selected for that variable over the full date range. The only exception to this is *rmbs*, which had a block size of 75 observations as shown in Table 5.5. Over the smaller date range, this particular choice of block size actually led to an inability to calculate a standard error for one of the parameters. So, the closest block size to 75 observations that did produce a tractable result (77 observations) was chosen for this variable on the common date range.

The GPD and GEVD parameter estimates for each variable are shown in Table 5.6. The general observations from estimating the marginal distributions over the full date ranges remain largely the same. We still see a wide variability in tail locations in accord with differences in the overall daily volatility of each asset, the tail shape parameters for certain assets are either statistically insignificant or only marginally significant at the typical levels, and the scale parameters are very statistically significant and generally in the 0.0015 – 0.0035 range.

Table 5.6. Marginal Distribution Estimation over Common Date Range

The table contains the marginal point estimates and observation counts for the includible life insurer asset classes. The GPD estimates are in Panel A, and the GEVD estimates are in Panel B. The *t*-statistics are included in parentheses. The number of extreme observations included in the tail or the block size chosen for each asset class is also included. Except for *port2013*, which is a portfolio of the individual asset classes using the 2013 industry-wide asset distribution as the weights, the number of tail observations modeled over this date range is equal to the same ratio of the number of total observations as in the full date range. Likewise, the block size used over this date range is set equal to the block size chosen in the full date range. The number of tail observations used for *port2013* was based on the Loretan and Phillips (1994) method and the block size was chosen from the same set used for the individual assets in the full date range.

Panel A – Generalized Pareto Distribution Estimates

	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*	*port2013*
$\hat{\lambda}$	-0.0060	-0.0064	-0.0040	-0.0144	-0.0020	-0.0031	-0.0001	-0.0001	-0.0061
$\hat{\tau}$	-0.0345	0.0056	0.0067	0.1802	0.0773	0.4415	0.4737	0.4087	0.3226
	(-0.73)	(0.12)	(0.17)	(3.13)	(1.93)	(8.61)	(4.09)	(5.36)	(4.30)
$\hat{\sigma}$	0.0033	0.0036	0.0024	0.0080	0.0014	0.0019	0.0001	0.0001	0.0023
	(21.31)	(20.10)	(36.28)	(13.77)	(679.32)	(946.68)	(38.62)	(50.14)	(19.71)
Tail Observations	404	454	404	404	420	496	110	220	248
Total Observations	4,236	4,236	4,236	4,236	4,236	4,234	4,236	4,236	4,234

Panel B – Generalized Extreme Value Distribution Estimates

	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*	*port2013*
$\hat{\lambda}$	-0.0072	-0.0044	-0.0048	-0.0152	-0.0037	-0.0036	-0.0001	-0.0001	-0.0025
	(-33.81)	(-34.07)	(-34.87)	(-23.88)	(-14.88)	(-22.03)	(-32.26)	(-74.62)	(-30.87)
$\hat{\tau}$	0.0819	0.0014	0.1419	0.2045	0.0376	0.3028	0.5273	0.7713	0.1237
	(1.34)	(0.06)	(2.07)	(3.39)	(0.41)	(4.14)	(4.57)	(4.79)	(4.96)
$\hat{\sigma}$	0.0029	0.0034	0.0020	0.0079	0.0018	0.0026	0.0001	0.0001	0.0024
	(63.37)	(67.14)	(987.25)	(17.37)	(908.84)	(1280.92)	(37.61)	(58.77)	(1204.08)
Block Size	21	5	21	21	77	21	126	126	5
Total Observations	4,236	4,236	4,236	4,236	4,236	4,234	4,236	4,236	4,234

Table 5.6 also includes the GPD and GEVD marginal distribution estimations for a portfolio of these asset classes (*port2013*). The weights used to construct this portfolio are based on the 2013 industry-wide distribution of combined assets (i.e., General and Separate Accounts combined). Due to the two missing observations for the *cmbs* variable, the portfolio also has only 4,234 observations on the common date range. The GPD estimation for the portfolio was done for the same four Loretan and Phillips cases used in the full date range. Likewise, the GEVD estimation was done under the same set of block sizes as the other variables. Although

some of the individual asset classes have statistically insignificant estimates for τ, the portfolio estimation produces a shape estimate with a p-value that is less than 1%.

G. Estimation of Joint Distribution

G. 1. Risk Hyperplane

Our initial attempt to model the joint tail risk of the includible asset classes is to make use of the concept of a risk hyperplane. To do this, we will estimate the Sortino Ratio across a variety of portfolio weights for a representative life insurance company's General Account. This analysis will focus on the company's choice of how much to invest in equities, which is an asset class with an especially high amount of market risk, and how much to invest in corporate bonds, which is the largest allocation in the General Account. Our representative life insurance company is one that matches the industry-wide asset allocations found in the 2014 ACLI Fact Book.

Although the General Account weights are the choice variables for the company, the Separate Account weights as of 2013 are included as an exogenous variable. It is appropriate to include the Separate Account weights exogenously because these reflect investment decisions made directly by policyholders rather than the company itself. Admittedly, the company retains some control because policyholders must choose from the investment options provided by the company. However, the ultimate decision of how much to invest in stocks, bonds, etc. lies with the policyholder. The degree to which the policyholders invest in asset classes with high market risk should also affect the company's own investment decisions in the General Account.

The weights actually used in the analysis are somewhat different than those derived directly from the ACLI Fact Books. This is because we do not have data on all types of life insurer assets. In particular, we lack data for assets such as policy loans, receivables, farm mortgages, and direct ownership of real estate. The weights for the includible asset classes are normalized to sum up to one after excluding those asset classes for which we lack data. From

2009 - 2013, the includible asset classes accounted for about 85% of the General Account assets with policy loans and miscellaneous assets, which includes premiums and investment income earned but not yet received by year-end, accounting for most of the gap.

As of 2013, the industry-wide General Account allocation to common stocks is 2.37% of the includible assets. This has been remarkably consistent in the post-financial crisis period, with the allocation ranging from 2.29% to 2.46% since the crisis. Going back twenty years to the mid-1990s, though, this allocation has been trending down. In fact, the 2013 allocation is 53% lower than that of 1994, when about 5% of the includible General Account assets were invested in equities. For our analysis, we will build twenty-one life insurer asset portfolios where the equity allocation has a low of 1% and increases in steps of 20 basis points each until we reach a maximum of 5.00%. The weights of the remaining includible assets will be based on their 2013 allocations. For example, corporate bonds make up 58.36% of the non-equity includible General Account assets as of 2013. If the portfolio under consideration includes an equity allocation of 1%, then the corporate bond allocation in this portfolio will be $0.5836 \times (1 - 0.01) \approx 0.5778$.

To estimate the tail risk and Sortino Ratios, we need to estimate the GPD and GEVD models for each of the twenty-one portfolios. For the GPD model, we will use a threshold corresponding to 248 tail observations, which is the number of tail observations selected for the portfolio using actual 2013 weights (*port2013* in Table 5.6). For the GEVD model, we will likewise use a block size of five to correspond to the selection made for *port2013*. The parameter estimates of these models will then be used to calculate the lower partial moment and Sortino Ratio for each portfolio.

The GPD and GEVD parameter estimates are provided in Table 5.7. Interestingly, the GPD model does not appear to be very sensitive to the choice of equity weight, at least within the range of equity weights considered. However, we can still see that adjusting the company's equity weight from 1% to 5% does have some impact on the tail risk of a company's portfolio. The location of the tail shifts from -0.0060 to -0.0064 and the scale increases from 0.0022 to 0.0023. Although these do not seem like big shifts in tail risk, they both move in the direction that greater allocations to equities increase the company's tail risk. It appears that the company does receive some compensation for this additional risk as the mean portfolio daily return increases from 0.0242 to 0.0243.

Table 5.7. Portfolio Marginal Distribution Estimation over Common Date Range

The table contains the GPD and GEVD parameter estimates and mean daily portfolio return for a range of equity portfolio weights. Panel A contains the estimates for the GPD model and Panel B contains the GEVD estimates. The equity weight in the portfolio under consideration is captured by the column headings within each panel. For every portfolio, the GPD model is estimated based on the 248 most extreme daily returns over the common date range and the GEVD model is estimated with a block size of five observations. The *t*-statistics are included in parentheses.

Panel A – Generalized Pareto Distribution					
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*
$\hat{\lambda}$	-0.0060	-0.0061	-0.0061	-0.0063	-0.0064
$\hat{\tau}$	0.3318	0.3244	0.3146	0.3352	0.3408
	(4.43)	(4.33)	(4.21)	(4.40)	(4.43)
$\hat{\sigma}$	0.0022	0.0023	0.0024	0.0023	0.0023
	(21.48)	(20.10)	(18.87)	(19.33)	(18.97)
Mean Return (× 100)	0.0242	0.0242	0.0243	0.0243	0.0243
Panel B – Generalized Extreme Value Distribution					
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*
$\hat{\lambda}$	-0.0024	-0.0024	-0.0025	-0.0025	-0.0026
	(-30.77)	(-30.85)	(-30.92)	(-30.99)	(-31.05)
$\hat{\tau}$	0.1219	0.1233	0.1247	0.1264	0.1277
	(4.90)	(4.95)	(4.99)	(5.04)	(5.09)
$\hat{\sigma}$	0.0024	0.0024	0.0024	0.0025	0.0025
	(1181.47)	(1197.85)	(1214.77)	(1231.35)	(1248.77)
Mean Return (× 100)	0.0242	0.0242	0.0243	0.0243	0.0243

To calculate the lower partial moment, we need to choose a target rate. Inspired by the Loretan and Phillips (1994) method for selecting the GPD threshold, the target rates for the portfolio lower partial moment estimations will be based on the 31, 62, 124, and 248 most extreme returns of the value-weighted stock index. These tail points correspond to daily stock returns of approximately -4.02%, -3.13%, -2.49%, and -1.86%, respectively. We then choose the target rate for the overall portfolio to be the minimum of the threshold from that portfolio's GPD estimation and the product of the stock return tail point and the equity portfolio weight (e.g., if the equity allocation is 1% and the equity market tail includes 31 observations, then the daily target rate return equals $0.01 \times -0.0402 = -0.000402$). We use the GPD threshold as a maximum to ensure that the target rate is located in the left tail used to estimate the GPD model. Based on the 2013 portfolio weights, the average daily portfolio return is equal to 0.000243.

After setting the target rate and equity weight for a given portfolio, we calculate the lower partial moment as described in Section C of this chapter. Table 5.8 contains the square root of the lower partial moment estimates for both the GPD and GEVD models. Recall that all of these estimates use the 248 most extreme portfolio daily returns for the GPD and a block size of five for the GEVD. So, the tail of the portfolio series used is the same for all estimates. Generally, this measure of tail risk produces lower estimates with the GEVD model than with the GPD model. However, the general conclusions are similar across both models.

We can see that the riskiness of the tail depends on both the target rate and the "aggressiveness" of the company's asset allocation as measured by equity weight. As the company moves to a less negative target rate, the tail risk decreases. In fact, the reduction in the square root of the lower partial moment is in the range of 0.25% to 0.28% for the GPD and 0.17% to 0.20% for the GEVD, depending on how much is invested in equities. The notable aspect of this result is that it is almost a 50% reduction in the tail risk of the most extreme target rate. This occurs because a more negative target rate focuses the analysis on the tip of the tail, which is the riskiest portion. As a greater portion of the portfolio's tail is included in the tail risk measure, more moderate manifestations of tail risk are included, which reduces the overall tail risk measure. Not surprisingly, investing more of the General Account in equities leads to some increase in tail risk. However, the sensitivity to equity weight is not as great as it is to the target rate. Increasing one's equity weight from 1% to 5% results in a tail risk increase of about two and a half to three basis points on average or as much as four to five basis points with a more extreme target rate.

Table 5.8. Portfolio Lower Partial Moment by Target Rate and Equity Weight

The table contains the square root of the lower partial moment estimates for portfolios of the includible asset classes. The portfolio composition is defined by the equity weights given in the column headings. The target rate used to calculate the lower partial moment is defined by the number of extreme observations used from the *vwst* series. The estimates are given in percentage terms (e.g., the square root of the lower partial moment for an equity weight of 1.00% and an equity tail of 31 observations is 0.4857%). The "High – Low" column calculates the difference between the estimate for a high equity weight of 5% and a low equity weight of 1%. The "Small Tail – Big Tail" row calculates the difference between the estimate for a smaller equity market tail of 31 observations and a bigger equity market tail of 248 observations. The GPD estimates (based on a portfolio tail of 248 observations) are in Panel A, and the GEVD estimates (based on a block size of five) are in Panel B.

Panel A – Generalized Pareto Distribution						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	High – Low
Equity Market Tail of 31 Obs.	0.4857	0.4890	0.5123	0.5245	0.5323	0.0466
Equity Market Tail of 62 Obs.	0.3669	0.3698	0.3721	0.3814	0.3875	0.0206
Equity Market Tail of 124 Obs.	0.2987	0.3060	0.3076	0.3229	0.3275	0.0288
Equity Market Tail of 248 Obs.	0.2339	0.2369	0.2427	0.2481	0.2546	0.0207
Small Tail – Big Tail	0.2518	0.2521	0.2696	0.2764	0.2777	
Panel B – Generalized Extreme Value Distribution						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	High – Low
Equity Market Tail of 31 Obs.	0.3712	0.3754	0.4001	0.4054	0.4102	0.0390
Equity Market Tail of 62 Obs.	0.2807	0.2844	0.2880	0.2922	0.2962	0.0155
Equity Market Tail of 124 Obs.	0.2377	0.2447	0.2471	0.2578	0.2608	0.0231
Equity Market Tail of 248 Obs.	0.1998	0.2029	0.2087	0.2123	0.2176	0.0178
Small Tail – Big Tail	0.1714	0.1725	0.1914	0.1931	0.1926	

Although it is clear from Tables 5.7 and 5.8 that higher allocations to equities result in a somewhat higher tail risk exposure for a life insurance company, it remains to be seen if they receive sufficient compensation for this additional exposure. When studying the risk-return trade-off from a downside risk perspective, the typical risk premium (i.e., the return from the risky asset net of the risk-free rate) is no longer appropriate. Instead, we calculate the portfolio's return in excess of the target rate. Obviously, if we set our target rate, or sometimes called the minimum acceptable rate, to be equal to the risk-free rate, then the excess return under the downside risk framework would be identical to the typical risk premium. However, sometimes it makes more sense to have a target rate different from the risk-free rate. For example, a life insurance company might set a portfolio target rate based on the rate of return that is required to maintain sufficient risk-based capital levels.

From Table 5.9, we can see that our representative life insurance does receive some compensation for taking on the extra tail risk exposure through either higher equity weights or accepting a more negative target rate. Excess returns are defined to be the average portfolio return for a given equity weight net of the target rate used. These excess returns are always at least 70 basis points higher when one focuses on a smaller portion of the tail than when one uses a larger portion of the tail. Excess returns are also higher as one increases the portfolio's allocation to equities. However, the amount of extra compensation received depends on the target rate. Average portfolio returns are nearly ten basis points higher when the equity weight is 5% relative to a weight of 1% when the most negative target rate is used. The amount of this extra return declines as one increases the target rate until only about five basis points of extra return is received with the least negative target rate. Still, the company can still expect to receive more return with a higher exposure to tail risk. To answer the question of whether or not this extra return is worth the extra risk, we need to review the Sortino Ratios of these portfolios.

Table 5.9. Portfolio Excess Return by Target Rate and Equity Weight

The table contains the average portfolio return in excess of the target rate for portfolios of the includible asset classes. The portfolio composition is defined by the equity weights given in the column headings, and the target rate is defined by the number of extreme observations used from the *vwst* series. The estimates are given in percentage terms (e.g., the excess return for an equity weight of 1.00% and an equity tail of 31 observations is 1.3494%). The "High – Low" column calculates the difference between the estimate for a high equity weight of 5% and a low equity weight of 1%. The "Small Tail – Big Tail" row calculates the difference between the estimate for a smaller equity market tail of 31 observations and a bigger equity market tail of 248 observations. Only one set of estimates is given since the portfolio excess returns are equivalent under the GPD and GEVD models.

Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	High – Low
Equity Market Tail of 31 Obs.	1.3494	1.3743	1.3992	1.4240	1.4489	0.0995
Equity Market Tail of 62 Obs.	1.0563	1.0757	1.0950	1.1144	1.1338	0.0775
Equity Market Tail of 124 Obs.	0.8441	0.8595	0.8749	0.8903	0.9056	0.0615
Equity Market Tail of 248 Obs.	0.6378	0.6494	0.6609	0.6724	0.6839	0.0461
Small Tail – Big Tail	0.7116	0.7249	0.7383	0.7516	0.7650	

Table 5.10 contains the Sortino Ratio estimates for these portfolios. Although there is some variation across particular combinations of the target rate and probability model used, some general conclusions are apparent. It is clear that life insurance companies should generally avoid being both conservative and aggressive with their equity allocations, as measured by the Sortino Ratio. The maximum Sortino Ratio for each target rate is in the "Optimal" column of Table 5.10, and this peak often occurs in the middle of the equity weights considered here (i.e., in the 1.5-3.0% range). In fact, averaging the equity weights corresponding to each maximal Sortino Ratio gives an optimal equity weight of exactly 2.00% for the GPD model and 2.50% for the GEVD model. Notably, these are quite close to the actual industry-wide equity allocation of 2.37% of includible assets as of 2013. Thus, we cannot reject the hypothesis that life insurers are making optimal asset allocation decisions in their General Account, given the allocation decisions of their policyholders in the Separate Account, from this analysis. Given the strong incentives to appropriately manage tail risk and the fact that life insurers are very informed investors, this result is not particularly surprising.

Table 5.10. Portfolio Sortino Ratio by Target Rate and Equity Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. The portfolio composition is defined by the equity weights given in the column headings, and the target rate is defined by the number of extreme observations used from the *vwst* series. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The maximum Sortino Ratio for each target rate and probability model is in the "Optimal" column with the corresponding equity weight in parentheses. The GPD estimates are in Panel A, and the GEVD estimates are in Panel B.

Panel A – Generalized Pareto Distribution						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	2.7783	2.8106	2.7310	2.7152	2.7222	2.8106 (2.00%)
Equity Market Tail of 62 Obs.	2.8790	2.9089	2.9428	2.9222	2.9258	2.9433 (3.20%)
Equity Market Tail of 124 Obs.	2.8255	2.8092	2.8441	2.7572	2.7650	2.8472 (1.40%)
Equity Market Tail of 248 Obs.	2.7273	2.7408	2.7226	2.7100	2.6862	2.7519 (1.40%)
Panel B – Generalized Extreme Value Distribution						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.6354	3.6609	3.4967	3.5127	3.5326	3.6609 (2.00%)
Equity Market Tail of 62 Obs.	3.7632	3.7825	3.8016	3.8140	3.8278	3.8278 (5.00%)
Equity Market Tail of 124 Obs.	3.5504	3.5124	3.5406	3.4536	3.4724	3.5711 (1.60%)
Equity Market Tail of 248 Obs.	3.1917	3.1997	3.1666	3.1679	3.1428	3.2143 (1.40%)

Our subsequent risk hyperplane analysis will proceed as follows. Next, we will study how sensitive the asset allocation optimality result is to a particular time period's allocations. This will be done by re-calculating the lower partial moments, excess returns, and Sortino Ratios for the same set of portfolios and target rates but with weights from different time periods. Recall that the analysis above was conducted given the Separate Account weights as of 2013. In addition, we also held fixed the proportion of the non-equity includible General Account allocations accounted for by each of the non-equity asset classes, which were also from 2013. Now, we will repeat the above analysis but using weights from 1994, 1998, 2002, 2006, and 2010.

Table 5.11. GPD-Based Sortino Ratios by Reference Year and Equity Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. Within each panel, the portfolio composition is defined by the equity weights given in the column headings, and the target rate is defined by the number of extreme observations used from the *vwst* series. However, the portfolio weights of the Separate Account assets and the non-equity General Account assets are based on the actual industry-wide allocations as of the year corresponding to each panel. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The maximum Sortino Ratio for each reference year and target rate is in the "Optimal" column, and the corresponding equity weight is in parentheses. The probability distribution function of the portfolio is modeled using a GPD model with a left tail of 248 observations.

Panel A – 2010						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	2.7119	2.8429	2.8622	2.7993	2.8209	2.8764 (3.20%)
Equity Market Tail of 62 Obs.	2.8587	2.8300	2.8461	2.8392	2.8603	2.8929 (1.80%)
Equity Market Tail of 124 Obs.	2.7931	2.8493	2.8507	2.8210	2.8480	2.8753 (2.80%)
Equity Market Tail of 248 Obs.	2.6904	2.7080	2.7106	2.7426	2.7429	2.7599 (4.20%)
Panel B – 2006						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	2.7725	2.7051	2.7311	2.6240	2.6167	2.8049 (1.80%)
Equity Market Tail of 62 Obs.	2.8835	2.9168	2.8915	2.8885	2.8764	2.9347 (2.20%)
Equity Market Tail of 124 Obs.	2.8017	2.7657	2.7956	2.7799	2.7028	2.8091 (1.40%)
Equity Market Tail of 248 Obs.	2.6929	2.7102	2.7012	2.6987	2.6494	2.7276 (2.40%)
Panel C – 2002						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	2.9937	2.8608	2.7377	2.7417	2.7354	2.9937 (1.00%)
Equity Market Tail of 62 Obs.	3.2071	3.1277	2.9895	2.9350	2.7878	3.2071 (1.00%)
Equity Market Tail of 124 Obs.	2.8586	2.8171	2.8299	2.7194	2.7181	2.8586 (1.00%)
Equity Market Tail of 248 Obs.	2.5661	2.6423	2.6476	2.6710	2.6643	2.6750 (4.20%)
Panel D – 1998						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	2.7697	2.7890	2.7454	2.6561	2.6196	2.7890 (2.00%)
Equity Market Tail of 62 Obs.	2.8484	2.8257	2.9262	2.8007	2.8192	2.9262 (3.00%)
Equity Market Tail of 124 Obs.	2.7167	2.7109	2.7733	2.7298	2.7773	2.8068 (4.60%)
Equity Market Tail of 248 Obs.	2.6660	2.6826	2.6798	2.6972	2.6939	2.6988 (4.60%)
Panel E – 1994						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.1774	3.2344	3.3996	3.4819	3.4524	3.5386 (4.40%)
Equity Market Tail of 62 Obs.	2.7544	2.9004	2.9330	2.9644	2.9095	3.0531 (3.60%)
Equity Market Tail of 124 Obs.	2.6948	2.6901	2.6926	2.6398	2.6319	2.7229 (2.60%)
Equity Market Tail of 248 Obs.	2.6948	2.6901	2.6926	2.6398	2.6271	2.7229 (2.60%)

Table 5.12. GEVD-Based Sortino Ratios by Reference Year and Equity Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. Within each panel, the portfolio composition is defined by the equity weights given in the column headings, and the target rate is defined by the number of extreme observations used from the *vwst* series. However, the portfolio weights of the Separate Account assets and the non-equity General Account assets are based on the actual industry-wide allocations as of the year corresponding to each panel. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The maximum Sortino Ratio for each reference year and target rate is in the "Optimal" column, and the corresponding equity weight is in parentheses. The probability distribution function of the portfolio is modeled using a GEVD model with a block size of five observations.

Panel A – 2010						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.5269	3.7512	3.7976	3.6525	3.6794	3.8069 (3.20%)
Equity Market Tail of 62 Obs.	3.7261	3.6888	3.7266	3.6807	3.7064	3.7564 (1.80%)
Equity Market Tail of 124 Obs.	3.4817	3.5707	3.5877	3.5239	3.5605	3.6124 (2.80%)
Equity Market Tail of 248 Obs.	3.1346	3.1622	3.1731	3.2007	3.2005	3.2191 (4.20%)
Panel B – 2006						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.6685	3.5287	3.5541	3.3964	3.4183	3.6988 (1.80%)
Equity Market Tail of 62 Obs.	3.8012	3.8251	3.7708	3.7886	3.8089	3.8301 (2.20%)
Equity Market Tail of 124 Obs.	3.5281	3.4637	3.4987	3.4977	3.4205	3.5441 (1.40%)
Equity Market Tail of 248 Obs.	3.1578	3.1716	3.1567	3.1642	3.1198	3.1823 (2.40%)
Panel C – 2002						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.9069	3.8845	3.6931	3.7855	3.7213	3.9990 (1.80%)
Equity Market Tail of 62 Obs.	4.0598	4.1355	3.9693	3.9874	3.7221	4.1529 (2.40%)
Equity Market Tail of 124 Obs.	3.4092	3.4628	3.5085	3.4258	3.4076	3.5236 (2.40%)
Equity Market Tail of 248 Obs.	2.8877	2.9258	2.9563	2.9757	3.0719	3.0719 (5.00%)
Panel D – 1998						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.7302	3.7976	3.6959	3.5877	3.4691	3.7976 (2.00%)
Equity Market Tail of 62 Obs.	3.7708	3.7723	3.9108	3.7637	3.7395	3.9108 (3.00%)
Equity Market Tail of 124 Obs.	3.3845	3.4074	3.4884	3.4680	3.5041	3.5427 (4.60%)
Equity Market Tail of 248 Obs.	3.1219	3.1193	3.1678	3.1988	3.1661	3.2002 (3.80%)
Panel E – 1994						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
Equity Market Tail of 31 Obs.	3.6483	3.7746	4.0609	4.1782	4.2101	4.3106 (4.60%)
Equity Market Tail of 62 Obs.	3.0762	3.2878	3.3622	3.4032	3.3590	3.5044 (3.60%)
Equity Market Tail of 124 Obs.	2.5383	2.6424	2.7904	2.8577	2.9568	2.9581 (4.80%)
Equity Market Tail of 248 Obs.	1.8957	1.9904	2.0804	2.1701	2.2608	2.2608 (5.00%)

For the sake of brevity, we are only presenting the Sortino Ratios of the various portfolios based on the weights from the historical reference years[10]. The GPD-based results are found in Table 5.11, and the GEVD-based results are found in Table 5.12. As before, the highest Sortino Ratio, along with the corresponding equity weight, for each reference year and target rate is in the "Optimal" column of each table. Note that the results for two different target rates using the GPD model and 1994 allocations are nearly identical. This occurs because the product of the 124th and 248th most negative equity returns and the Combined Account equity allocation is closer to zero than the threshold used in the portfolio's GPD model. The target rate of the portfolio must be at least as far from zero in absolute value as the threshold used in the portfolio's GPD or GEVD modeling. When the target rate is calculated to be between zero and the portfolio's threshold, then the target rate is set equal to the threshold. As a result, two of the equity market tail sizes in 1994 are using the same target rate for nearly all of the potential equity weights.

The actual General Account equity allocations for the industry in these years are 2.46% (2010), 3.07% (2006), 2.94% (2002), 4.54% (1998), and 5.04% (1994). In other words, life insurance companies have progressively reduced the amount they choose to invest in equities over the past twenty years. However, our estimates of the optimal equity allocations do not match the actual investment behavior of the industry quite as well as in 2013. Figure 5.1 helps us visualize a comparison of the actual General Account equity allocations with the optimal allocations, averaged across target rate, as determined by our risk hyperplane analysis.

[10] Data on the probability model estimation, portfolio lower partial moments, and portfolio excess returns are available upon request.

Figure 5.2. Optimal and Actual General Account Equity Allocations (Risk Hyperplane)
This graph plots the optimal General Account equity allocations as determined by the GPD and GEVD-based portfolio models as compared with the historical industry-wide allocations. These allocations are based on historical portfolio weights for Separate Account assets and non-equity General Account assets from various years, which are shown on the horizontal axis. The vertical axis denotes the percentage of the General Account allocated to equities.

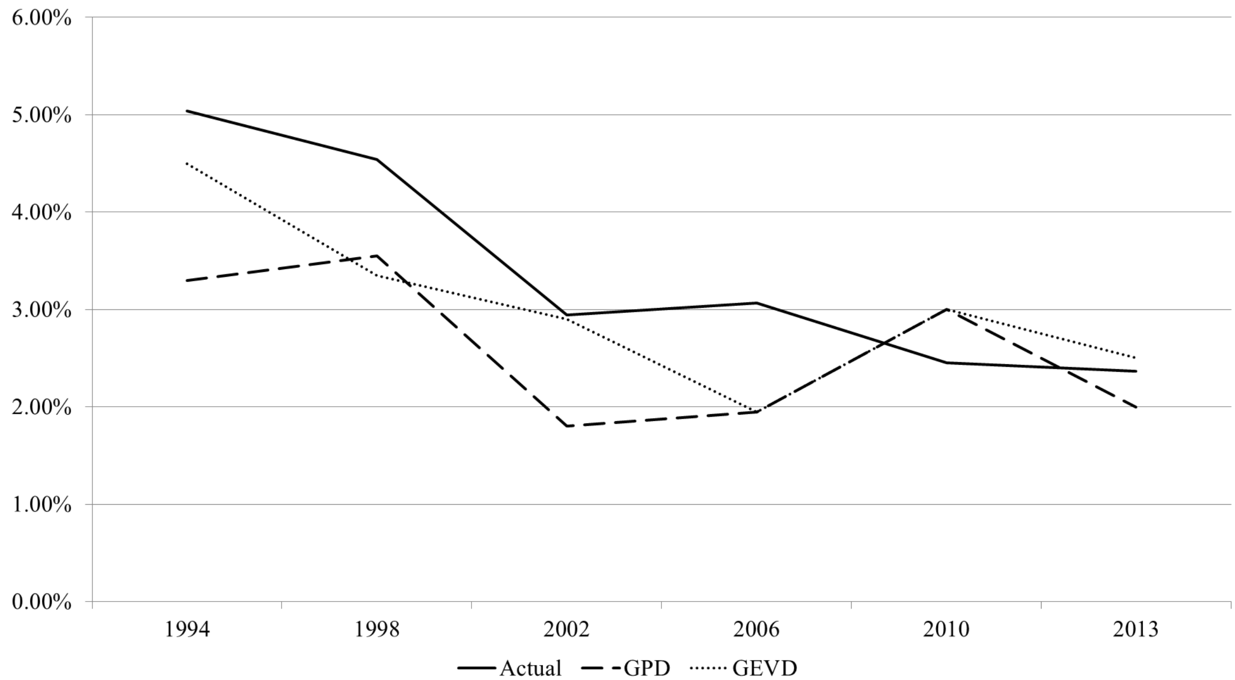

In the earlier years of the analysis (i.e., from 1994 to 2006), our models suggest that the industry was over-weighting equities relative to an "optimal" level, although the degree of the overweighting depends on whether or not we use the GPD or GEVD model. Using the GEVD model, the overweighting was relatively minor (about 0.72% on average), but it was more significant (about 1.25% on average) compared to the optimal allocations produced by the GPD model. Noticeably, both the actual and optimal equity levels are decreasing throughout this period, so there is some consistency between the changes in actual and optimal allocations even if the levels are different. Perhaps this is being driven by an increase in the equity exposure from the Separate Account. In fact, both the Separate Account's equity allocation and the overall Separate Account weight increased by about eighteen percentage points each (from 66.64% to

84.80% and from 18.03% to 35.78%, respectively). Beyond 2006, the trend changes and the actual and optimal equity levels become more stable. At this point, both models produce optimal allocations that are quite consistent with each other once we average across target rate. In 2010, the models actually suggest that the industry has now started underweight equities in their General Accounts, but this seems to have been corrected by 2013. As we noted earlier, the industry and our models are quite closely aligned at the end of our analysis in 2013.

Although there are hints from Figure 5.1 that our models may have some contrarian aspects (underweight equities during the market run-ups of the late 1990s and prior to the financial crisis, the pace of the downward trend in equity allocations slows for the GEVD following the 2001 market crash, and overweight equities after the financial crisis), we are not ready to claim our models have such prescience. Recall, that the only thing we are changing as we move from one reference year to the next are the portfolio weights of assets other than General Account equities. We are still estimating the probability models, risk measures, and portfolio returns using the full time series of asset returns. So, it appears that the fluctuation in the optimal allocations to General Account equities may be better explained by noticing how the make-up of the other assets changes. However, there are a number of allocation changes in the other assets from reference year to reference year and each can have an impact on the portfolio's tail risk, returns, and optimal allocations. Thus, to properly analyze this particular question requires a more robust multivariate analysis than is the purpose of the current study. Also, the low number of observations limits our ability to make strong conclusions based on changes in the optimal allocations from one reference year to the next.

Figure 5.3. Combined Account Tail Risk by Reference Year
This graph plots the average tail risk, as measured by the square root of the second lower partial moment, for each asset allocation reference year. Given the twenty-one possible General Account equity weights and the four target rates used, the value plotted in the graph for each probability model is the average value across all eighty-four of these equity weight-target rate combinations. Each reference year is shown on the horizontal axis, and the vertical axis denotes the average tail risk measurement in percentage points.

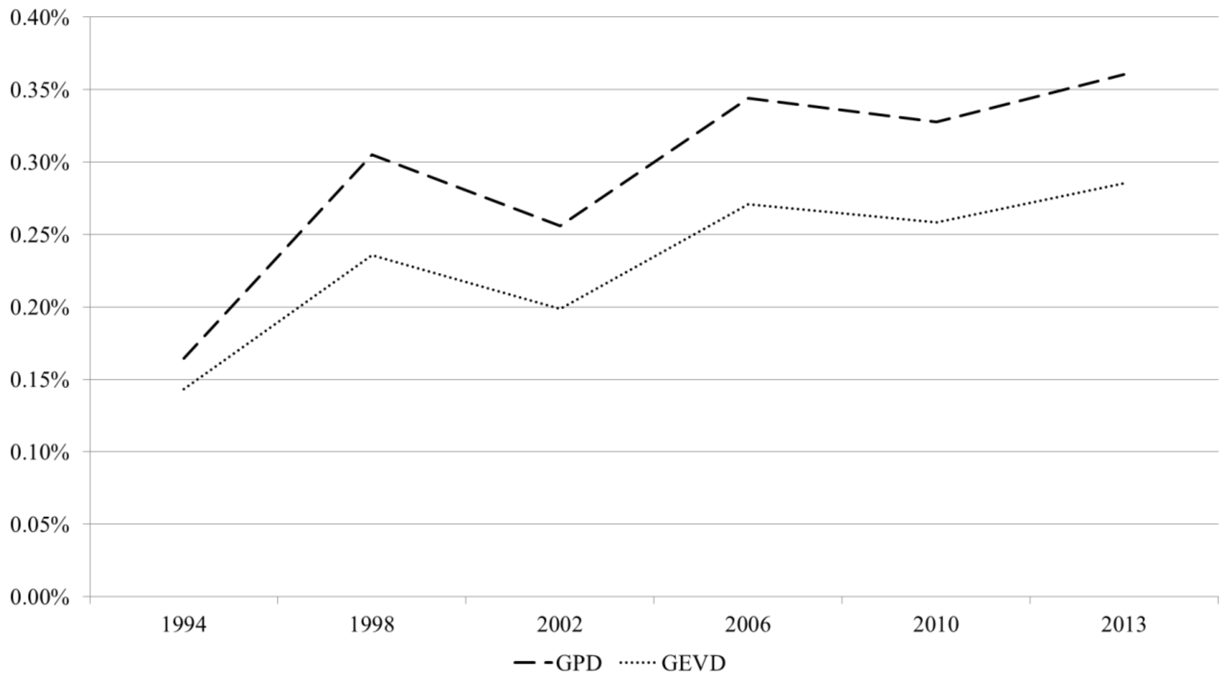

Still, because we know that all of the fluctuation in the optimal allocations must be due to changes in the exogenous allocations, this may help suggest some broad trends. The most significant shift that occurs in these exogenous allocations over this time period is the large increase in equity exposure coming from the Separate Account, which we just referred to a little bit ago. It is also clear that the tail risk of the Combined Account has significantly increased over this same time period, which is likely due in large part to the increased Separate Account equity exposure. We know this by looking at the tail risk of the portfolios under consideration by our representative life insurer across the various reference years. Over the course of the time period under analysis, the square root of the second lower partial moment of the portfolios under consideration, when averaged across the target rate used and the equity weight chosen, increases by about 99% to 119% depending on the probability model used (see Figure 5.3 for a graphical illustration). For sure, most of this increase occurs between 1994 and 1998 when the relative size of the Separate Account and the equity allocation within the Separate Account both jump to higher levels, but there is still a non-trivial increase since then. Naturally, it makes sense for the life insurers to respond to this significant increase in tail risk by pulling back somewhat on their own equity allocations, which is what we see in the optimal allocations and in the actual behavior of the industry over this time period.

Next, we will conduct a similar risk hyperplane analysis but with an alternative focus. Instead of focusing on equities, which is a small General Account allocation but a primary exposure to tail risk in an insurer's Combined Account, we will now shift our focus to corporate bonds. The corporate bond allocation is important for a life insurer because it is, by far, the largest allocation directly under the company's control. As of 2013, this asset class comprised nearly 57% of the includible General Account assets, which was just under four times larger than the second highest General Account allocation. The trend for corporate bonds has been the opposite of equities since 1994. Its allocation has increased nearly eight percentage points (from about 49% in 1994 to about 57% in 2013) over the past twenty years. One hypothesis, then, is that life insurers have responded to the increased tail risk exposure from higher Separate Account equity allocations by shifting more of their own investments into corporate bonds. This allows them to reduce the tail risk exposure of the General Account without sacrificing as much potential return as shifting into government bonds.

As with the equities, we analyze portfolios of includible life insurer assets that vary based on their General Account corporate bond allocation. Initially, the corporate bond allocation was to vary from a low of 48% to a high of 58% to reflect the range of historical General Account weights in this asset class. However, the initial results showed that the maximum weight was a binding constraint in almost all cases. So, we looked at some alternative ranges to better capture the maximal Sortino Ratio points. We found that our risk hyperplane analysis apparently shares a similar downside as the classical Markowitz portfolio optimization approach. Recall that one issue with the classical Markowitz approach is that it can produce allocations for particular assets that are much higher than might otherwise be reasonable, especially for assets that have a relatively high historical Sharpe Ratio. We run into a similar issue here with corporate bonds. From Table 5.2, we see that corporate bonds have an average daily return only slightly smaller than equities and higher than many of the other asset classes. However, it also has one of the smallest daily variances, with the exception of the two U.S. Treasury bills series, so it may very well have relatively lower tail risk. As a result, our risk hyperplane approach wants to select General Account allocations for corporate bonds that are in the 80% to 100% range. Clearly, these allocations may not be reasonable for the typical life insurance company given that the industry has not invested anywhere close to that much in corporate bonds in the past. Nonetheless, we believe some conclusions can still be obtained by analyzing the optimal allocations in an important asset class for life insurers.

For all of the portfolios, the Separate Account allocations to all asset classes and the General Account allocations to non-corporate bond assets are held fixed. We estimate the probability distributions of these portfolios using the GPD and GEVD models and estimate the lower partial moments, excess returns, and Sortino Ratios of the portfolios in order to select an optimal corporate bond allocation from a representative insurer's perspective. Using a range of 80% to 100% and a step size of 1% as our set of feasible corporate bond allocations, the Sortino Ratios of these portfolios are presented in Tables 5.13 and 5.14.

Table 5.13. GPD-Based Sortino Ratios by Reference Year and Corporate Bond Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. Within each panel, the portfolio composition is defined by the corporate bond weights given in the column headings, and the target rate is defined by the number of extreme observations used from the *corp* series. However, the portfolio weights of the Separate Account assets and the non-corporate bond General Account assets are based on the actual industry-wide allocations as of the year corresponding to each panel. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The maximum Sortino Ratio for each reference year and target rate is in the "Optimal" column, and the corresponding corporate bond weight is in parentheses. The probability distribution function of the portfolio is modeled using a GPD model with a left tail of 248 observations.

Panel A – 2013						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.6915	2.7121	2.7299	2.7567	2.8640	2.8640 (100%)
Corporate Bond Tail of 62 Obs.	2.6915	2.7121	2.7299	2.7382	2.7382	2.7458 (97%)
Corporate Bond Tail of 124 Obs.	2.6915	2.7121	2.7299	2.7382	2.7382	2.7458 (97%)
Corporate Bond Tail of 248 Obs.	2.6915	2.7121	2.7299	2.7382	2.7382	2.7458 (97%)
Panel B – 2010						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.7602	2.7764	2.7644	2.8178	2.7958	2.8795 (97%)
Corporate Bond Tail of 62 Obs.	2.7602	2.7677	2.7686	2.7670	2.7348	2.7755 (88%)
Corporate Bond Tail of 124 Obs.	2.7602	2.7677	2.7686	2.7670	2.7348	2.7755 (88%)
Corporate Bond Tail of 248 Obs.	2.7602	2.7677	2.7686	2.7670	2.7348	2.7755 (88%)
Panel C – 2006						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.7146	2.7121	2.8308	2.7744	2.7657	2.8583 (91%)
Corporate Bond Tail of 62 Obs.	2.7146	2.7121	2.7369	2.7654	2.7753	2.7768 (98%)
Corporate Bond Tail of 124 Obs.	2.7146	2.7121	2.7369	2.7654	2.7753	2.7768 (98%)
Corporate Bond Tail of 248 Obs.	2.7146	2.7121	2.7369	2.7654	2.7753	2.7768 (98%)
Panel D – 2002						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.9012	3.0991	3.1616	3.2962	3.3771	3.3771 (100%)
Corporate Bond Tail of 62 Obs.	2.6069	2.7108	2.8383	2.8649	2.9424	2.9424 (100%)
Corporate Bond Tail of 124 Obs.	2.6733	2.6796	2.6363	2.6453	2.6957	2.6957 (100%)
Corporate Bond Tail of 248 Obs.	2.6733	2.6796	2.6363	2.6231	2.6348	2.6796 (85%)
Panel E – 1998						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.7164	2.7420	2.8446	2.8847	3.0024	3.0076 (99%)
Corporate Bond Tail of 62 Obs.	2.7150	2.7210	2.7091	2.7316	2.7151	2.7910 (94%)
Corporate Bond Tail of 124 Obs.	2.7150	2.7210	2.7095	2.6827	2.6740	2.7265 (87%)
Corporate Bond Tail of 248 Obs.	2.7150	2.7210	2.7095	2.6827	2.6740	2.7265 (87%)
Panel E – 1994						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	4.0824	4.2267	4.1974	4.1364	4.3372	4.4581 (99%)
Corporate Bond Tail of 62 Obs.	3.3815	3.6360	3.7007	3.8961	4.1625	4.1625 (100%)
Corporate Bond Tail of 124 Obs.	2.9942	2.9956	3.1111	3.2702	3.2591	3.3438 (98%)
Corporate Bond Tail of 248 Obs.	2.7895	2.7984	2.7887	2.8836	2.8842	2.8842 (100%)

Table 5.14. GEVD-Based Sortino Ratios by Reference Year and Corporate Bond Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. Within each panel, the portfolio composition is defined by the corporate bond weights given in the column headings, and the target rate is defined by the number of extreme observations used from the *corp* series. However, the portfolio weights of the Separate Account assets and the non-corporate bond General Account assets are based on the actual industry-wide allocations as of the year corresponding to each panel. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The maximum Sortino Ratio for each reference year and target rate is in the "Optimal" column, and the corresponding corporate bond weight is in parentheses. The probability distribution function of the portfolio is modeled using a GEVD model with a block size of five observations.

Panel A – 2013						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.6995	2.8848	3.0773	3.2165	3.3787	3.3787 (100%)
Corporate Bond Tail of 62 Obs.	2.3036	2.3773	2.5139	2.6455	2.8074	2.8074 (100%)
Corporate Bond Tail of 124 Obs.	1.7992	1.9290	2.0539	2.2141	2.2900	2.2900 (100%)
Corporate Bond Tail of 248 Obs.	1.3973	1.4813	1.5608	1.6394	1.7682	1.7682 (100%)
Panel B – 2010						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	3.0956	3.2859	3.3157	3.4348	3.4340	3.5193 (97%)
Corporate Bond Tail of 62 Obs.	2.4763	2.6960	2.8553	3.0031	3.2359	3.2359 (100%)
Corporate Bond Tail of 124 Obs.	2.0730	2.2004	2.3092	2.4848	2.5662	2.5662 (100%)
Corporate Bond Tail of 248 Obs.	1.5679	1.6981	1.7941	1.8886	2.0136	2.0136 (100%)
Panel C – 2006						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	2.9243	3.1158	3.3264	3.3287	3.3706	3.4549 (94%)
Corporate Bond Tail of 62 Obs.	2.4021	2.5153	2.7402	2.8962	3.0396	3.0396 (100%)
Corporate Bond Tail of 124 Obs.	1.9640	2.1051	2.2459	2.3394	2.4852	2.4852 (100%)
Corporate Bond Tail of 248 Obs.	1.4899	1.5930	1.7094	1.7993	1.8962	1.8962 (100%)
Panel D – 2002						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	3.1095	3.3532	3.3604	3.4692	3.5389	3.5389 (100%)
Corporate Bond Tail of 62 Obs.	2.7156	2.8472	2.9493	2.9595	3.0323	3.0323 (100%)
Corporate Bond Tail of 124 Obs.	2.7624	2.7121	2.6628	2.6435	2.7293	2.7658 (81%)
Corporate Bond Tail of 248 Obs.	2.7624	2.7121	2.6628	2.6435	2.6371	2.7658 (81%)
Panel E – 1998						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	3.2291	3.3083	3.4644	3.5147	3.7016	3.7048 (99%)
Corporate Bond Tail of 62 Obs.	2.8136	3.0397	3.1553	3.1946	3.1994	3.2720 (98%)
Corporate Bond Tail of 124 Obs.	2.3017	2.4362	2.6039	2.7049	2.8541	2.8541 (100%)
Corporate Bond Tail of 248 Obs.	1.7725	1.8854	1.9894	2.1212	2.2310	2.2310 (100%)
Panel E – 1994						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
Corporate Bond Tail of 31 Obs.	4.9191	5.1444	5.1048	5.1358	5.4427	5.5695 (99%)
Corporate Bond Tail of 62 Obs.	3.9117	4.2294	4.3155	4.6753	5.0834	5.0834 (100%)
Corporate Bond Tail of 124 Obs.	3.3927	3.3896	3.5326	3.7814	3.7913	3.8599 (98%)
Corporate Bond Tail of 248 Obs.	2.6034	2.7484	2.9573	3.0617	3.2466	3.2466 (100%)

Again, these optimal allocations are much higher than those seen in the actual investing choices of the life insurance industry. Other than any potential issues with estimation error or sensitivity in our risk hyperplane models, this may also be exhibiting another limitation of our study. We have chosen to focus on only one of the major sources of risk for life insurance companies. When discussing how much to invest in corporate bonds, another major risk that we exclude from the analysis is credit risk. Given that the life insurance industry already finances a significant part of the corporate bond market, greatly expanding the General Account's allocations to corporate bonds, as our models suggest here, may force life insurance companies to move into riskier classes of bonds, which would offset at least some of the market risk benefits of such a strategy.

G. 2. Copulas

To complete the analysis, we model the joint distribution of the includible asset classes by making use of copula theory. In particular, we will need to build a vine of bivariate copula pairs using the techniques of Czado (2010), Czado, Schepsmeier, and Min (2012), and Brechmann and Schepsmeier (2013). Instead of needing to draw conclusions about optimality based on empirical estimates as in the risk hyperplane analysis, copulas allow us to be more precise and theoretical in pinpointing the optimal asset allocation decisions for life insurers concerned about long-term solvency.

As described in Section V.D.2, we first model the marginal distribution of each includible asset using a two-tailed GPD model. Consistent with the marginal GPD estimation performed earlier, we set the threshold for each asset class's lower tail to be equal to that used earlier in Table 5.6. The upper tail's threshold is then set equal to the lower tail's threshold except with the opposite sign. Table 5.15 contains the estimated parameters of these GPD marginal distributions. Again, it is apparent that equities contain a significant amount of tail risk for life insurers. As a class, they are fairly heavy-tailed with above-average tail shape parameter estimates and the highest scale parameter estimates. Corporate bonds, on the other hand, exhibit much flatter tails and more reasonable dispersion as measured by the scale parameter.

Table 5.15. GPD Marginal Distribution Estimation for Vine Copulas

The table contains the GPD parameter estimates for each includible asset class. For each asset class, the lower tail threshold is based on the return associated with the corresponding tail observations in Panel A of Table 5.6. The upper tail threshold is then set equal to the lower tail threshold but with the opposite sign. The threshold is denoted by λ, the tail shape by τ, and the scale by σ. The subscript on each parameter indicates whether the estimate is for the lower or upper tail.

	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*
$\hat{\lambda}_l$	-0.0060	-0.0064	-0.0040	-0.0144	-0.0020	-0.0031	-0.0001	-0.0001
$\hat{\lambda}_u$	0.0060	0.0064	0.0040	0.0144	0.0020	0.0031	0.0001	0.0001
$\hat{\tau}_l$	-0.0109	0.0144	-0.0308	0.1797	0.1065	0.4587	0.4639	0.4197
$\hat{\tau}_u$	0.1063	0.0658	0.0087	0.2387	0.1365	0.3645	-0.0048	0.0707
$\hat{\sigma}_l$	0.0032	0.0035	0.0025	0.0079	0.0013	0.0019	0.0001	0.0001
$\hat{\sigma}_u$	0.0026	0.0033	0.0018	0.0071	0.0012	0.0020	0.0001	0.0002

After estimating the marginal distributions for each asset class using the GPD model, we use the corresponding cumulative distribution functions to produce data for our assets that can be used in vine copula. Recall that a copula function requires as inputs the marginal cumulative distribution functions. Before we can estimate the vine copula, though, we need to select an appropriate structure. For a C-vine copula, this means we must select the root variable for each tree of the vine. To do this, we utilize the Kendall's tau estimates of dependence between each pair of variables in the tree. The Kendall's tau estimates for the first tree are given in Table 5.16, and it becomes clear that corporate bonds are chosen to be the root variable of the first tree as it has a higher sum of tau estimates than any other asset class. These estimates are calculated anew among the remaining seven asset classes for the second tree in the vine except these tau estimates are now conditional on the corporate bond returns. Again, the asset class with the highest aggregate joint dependency, using the absolute value of each individual tau estimate, is chosen to be the root variable for the second tree. This process repeats itself until the root variables for all of the trees are chosen. The order of these root variables for our study is *corp*, *trbd6*, *rmbs*, *trbd*, *cmbs*, *fnbd*, and *trbd3* and *vwst* sharing equally in the final tree.

Table 5.16. C-Vine Copula Root Variable Selection

The table contains the Kendall's tau estimates for each pair of the includible life insurer asset classes. These are used to select the appropriate root variable for the first tree of a C-vine copula. The *Sum* column contains the aggregate sum over each row when using the absolute value of each Kendall's tau estimate in that row. The highest sum is in boldface and corresponds to the asset class selected to be the root variable for the first tree of the vine copula.

	trbd	*fnbd*	*corp*	*vwst*	*rmbs*	*cmbs*	*trbd3*	*trbd6*	*Sum*
trbd	1.0000	0.2218	0.8165	-0.1885	0.6243	0.5195	0.0740	0.2016	3.6461
fnbd	0.2218	1.0000	0.2201	-0.0492	0.2188	0.2029	0.0349	0.1003	2.0480
corp	0.8165	0.2201	1.0000	-0.1537	0.6413	0.5777	0.0632	0.1959	**3.6685**
vwst	-0.1885	-0.0492	-0.1537	1.0000	-0.1095	-0.0665	-0.0458	-0.0811	1.6942
rmbs	0.6243	0.2188	0.6413	-0.1095	1.0000	0.5528	0.0849	0.2254	3.4569
cmbs	0.5195	0.2029	0.5777	-0.0665	0.5528	1.0000	0.0629	0.1811	3.1635
trbd3	0.0740	0.0349	0.0632	-0.0458	0.0849	0.0629	1.0000	0.5738	1.9396
trbd6	0.2016	0.1003	0.1959	-0.0811	0.2254	0.1811	0.5738	1.0000	2.5593

For a D-vine copula, we make use of the Kendall's tau estimates again. However, this time, the basic structure of the vine is different. Instead of looking for the single variable with the maximum joint dependence across all of the asset classes, we are looking for the order that produces the maximum aggregate joint dependence. Recall that the bivariate pairs in the first tree of a D-vine copula are setup differently than in a C-vine copula. Instead of having each non-root variable paired with the root variable, which is the case in a C-vine copula, each variable is paired with only the variable immediately preceding and following it in the order (e.g., if the variable order is 1, 2, 3, and 4, then the pairs for the first tree would be 1-2, 2-3, and 3-4). With eight asset classes, we have 8!, or 40,320, possible orders for the first tree. Setting up all of the possible ordering schemes and using the Kendall's tau estimates for each pair from Table 5.16 allows us to find the order with the highest aggregate joint dependence. Again, we use the absolute value of each tau estimate when taking the sum across all of the pairs in the ordering scheme. This produces an optimal D-vine copula order of *trbd3*, *trbd6*, *fnbd*, *cmbs*, *rmbs*, *corp*, *trbd*, and *vwst*.

After choosing an appropriate vine copula structure, we must select an appropriate copula function for each bivariate pair in the vine and estimate the corresponding parameters of the chosen bivariate copula. Before reviewing the vine copula estimation results, we review some of the major bivariate copulas (see Czado, Schepsmeier, and Min (2012) and Brechmann and Schepsmeier (2013) for further details). In finance, some commonly used copula functions include the Gaussian and Student-*t* copulas, which both belong to a class called elliptical copulas. Both of these have symmetric tail dependence but in different ways. The Gaussian copula has zero tail dependence while it is non-zero for the Student-*t* copula. An alternative class of functions is that of the Archimedean copulas. These copulas allow for more flexible forms of tail dependence. For example, the Clayton copula exhibits non-zero lower tail dependence and zero upper tail dependence, the Gumbel and Joe copulas exhibit the opposite, and the Frank copula has zero tail dependence in both tails. These copulas are functions of a single dependence parameter where the higher degrees of dependence are associated with higher parameter values. They also have variants that are based on the underlying copula function but have alternative dependency structures. For example, the Survival Gumbel copula is based on the Gumbel copula function but with a non-zero lower (instead of upper) tail dependence and a zero upper (instead of lower) tail dependence. Some of the Archimedean copulas are even more flexible by allowing for dependencies that are potentially asymmetric and both non-zero. These copulas are governed by two parameters to allow for such flexibility. In light of this bivariate copula review, the copula selection and estimation results for our C-vine copula are presented in Table 5.17.

One notable observation of the vine copula estimation results is the frequency with which the Student-*t* copula is chosen. Given that the Student-*t* distribution is similar to the oft-used Gaussian distribution but with somewhat heavier tails, perhaps this is not very surprising. Nonetheless, it is interesting to see how often a relatively simple dependence structure fits asset return pairs better than more complicated dependence structures. Another observation is that the dependence within the pairs (as measured by *Parameter 1* in Table 5.17, which is in (-1, 1) for the Student-*t* copula and often greater than one for the Archimedean copulas) declines significantly after conditioning on the corporate bond returns. However, the estimated degrees of freedom, which is a factor in the tail dependence modeled by the copula, does seem to be generally higher after conditioning on the corporate bond returns. Still, this suggests that corporate bond returns are driving much of the joint dependencies of life insurance company investment portfolios making a C-vine copula structure a reasonable approach. This is also reflected in the fact that the independence copula is chosen more frequently in the later trees of the vine.

Table 5.17. C-Vine Copula Selection and Estimation

The table contains the copula selection and parameter estimates for each bivariate pair in the C-vine copula of a life insurer's includible asset classes. *Variable 1* contains the root variable for each respective tree in the vine. *Conditioning Set* contains the root variables from earlier trees that are now conditioned on when estimating the current tree. The entries in *Copula* indicate the copula function chosen where "I" denotes the independence copula, "N" denotes the Gaussian copula, "t" denotes the Student-*t* copula, "F" denotes the Frank copula, "BB8" denotes the Joe-Frank two-parameter Archimedean copula, "SG" denotes the survival version of the Gumbel copula, and "SJ" denotes the survival version of the Joe copula. For the Gaussian and Student-*t* copulas, *Parameter 1* is a dependence parameter $\rho \in (-1, 1)$. For the Student-*t* copula, *Parameter 2* is a degrees of freedom parameter $\nu > 2$. For the one-parameter Archimedean copulas, *Parameter 1* is $\theta \geq 1$ (for Gumbel), $\theta \in \mathbb{R} \setminus \{0\}$ (for Frank), or $\theta > 1$ (for Joe). For the two-parameter Archimedean copulas, the dependence is governed by two parameters, which are $\theta \geq 1$ and $\delta \in (0, 1]$ for the BB8 copula.

Tree	*Variable 1*	*Variable 2*	*Conditioning Set*	*Copula*	*Parameter 1*	*Parameter 2*
1	*corp*	*trbd6*	N/A	SG	1.2499	-
	corp	*rmbs*	N/A	t	0.8310	4.5747
	corp	*trbd*	N/A	t	0.9559	2.2150
	corp	*cmbs*	N/A	t	0.7978	2.1251
	corp	*fnbd*	N/A	t	0.3381	6.7571
	corp	*trbd3*	N/A	SJ	1.1079	-
	corp	*vwst*	N/A	t	-0.2431	4.6258
2	*trbd6*	*rmbs*	*corp*	t	0.1819	8.0417
	trbd6	*trbd*	*corp*	I	-	-
	trbd6	*cmbs*	*corp*	t	0.1307	10.2479
	trbd6	*fnbd*	*corp*	t	0.0685	15.8650
	trbd6	*trbd3*	*corp*	BB8	4.3970	0.9164
	trbd6	*vwst*	*corp*	t	-0.0661	6.8569
3	*rmbs*	*trbd*	*corp*, *trbd6*	I	-	-
	rmbs	*cmbs*	*corp*, *trbd6*	t	0.2987	5.8977
	rmbs	*fnbd*	*corp*, *trbd6*	t	0.1037	16.5928
	rmbs	*trbd3*	*corp*, *trbd6*	t	-0.0614	12.3659
	rmbs	*vwst*	*corp*, *trbd6*	t	0.0844	8.7589
4	*trbd*	*cmbs*	*corp*, *trbd6*, *rmbs*	t	-0.1859	7.3073
	trbd	*fnbd*	*corp*, *trbd6*, *rmbs*	I	-	-
	trbd	*trbd3*	*corp*, *trbd6*, *rmbs*	t	0.0260	26.1849
	trbd	*vwst*	*corp*, *trbd6*, *rmbs*	t	-0.2211	9.1631
5	*cmbs*	*fnbd*	*corp*, *trbd6*, *rmbs*, *trbd*	N	0.0722	-
	cmbs	*trbd3*	*corp*, *trbd6*, *rmbs*, *trbd*	t	-0.0603	15.8395
	cmbs	*vwst*	*corp*, *trbd6*, *rmbs*, *trbd*	I	-	-
6	*fnbd*	*trbd3*	*corp*, *trbd6*, *rmbs*, *trbd*, *cmbs*	F	-0.2615	-
	fnbd	*vwst*	*corp*, *trbd6*, *rmbs*, *trbd*, *cmbs*	I	-	-
7	*trbd3*	*vwst*	*corp*, *trbd6*, *rmbs*, *trbd*, *cmbs*, *fnbd*	F	0.2397	-

Table 5.18. D-Vine Copula Selection and Estimation

The table contains the copula selection and parameter estimates for each bivariate pair in the D-vine copula of a life insurer's includible asset classes. The entries in *Copula* indicate the copula function chosen where "I" denotes the independence copula, "t" denotes the Student-*t* copula, "F" denotes the Frank copula, "BB8" denotes the Joe-Frank two-parameter Archimedean copula, "SBB7" denotes the survival version of the Joe-Clayton two-parameter Archimedean copula, and "RBB8" denotes the 270-degree rotated version of the BB8 copula. For the Student-*t* copula, *Parameter 1* is a dependence parameter $\rho \in (-1, 1)$ and *Parameter 2* is a degrees of freedom parameter $\nu > 2$. For the Frank one-parameter Archimedean copula, *Parameter 1* is $\theta \in \mathbb{R} \setminus \{0\}$. For the two-parameter Archimedean copulas, the dependence is governed by two parameters, which are $\theta \geq 1$ and $\delta > 0$ for the BB7 copula and $\theta \geq 1$ and $\delta \in (0, 1]$ for the BB8 copula. The 270-degree rotated versions of the Archimedean copulas allow for negative dependence and have parameter spaces with the opposite sign (i.e., $\theta \leq -1$ and $\delta \in [-1, 0)$ for the RBB8 copula).

Tree	*Variable 1*	*Variable 2*	*Conditioning Set*	*Copula*	*Parameter 1*	*Parameter 2*
1	*trbd3*	*trbd6*	N/A	BB8	4.7491	0.8859
	trbd6	*fnbd*	N/A	SBB7	1.1160	0.1075
	fnbd	*cmbs*	N/A	t	0.3110	6.1904
	cmbs	*rmbs*	N/A	t	0.7587	2.8094
	rmbs	*corp*	N/A	t	0.8310	4.5747
	corp	*trbd*	N/A	t	0.9559	2.2150
	trbd	*vwst*	N/A	t	-0.2961	4.2864
2	*trbd3*	*fnbd*	*trbd6*	F	-0.6923	-
	trbd6	*cmbs*	*fnbd*	t	0.2565	4.9743
	fnbd	*rmbs*	*cmbs*	t	0.1557	18.4851
	cmbs	*corp*	*rmbs*	t	0.4699	4.1301
	rmbs	*trbd*	*corp*	I	-	-
	corp	*vwst*	*trbd*	t	0.1314	6.4029
3	*trbd3*	*cmbs*	*trbd6, fnbd*	RBB8	-1.7545	-0.7682
	trbd6	*rmbs*	*fnbd, cmbs*	t	0.1692	21.0446
	fnbd	*corp*	*cmbs, rmbs*	t	0.0668	12.2309
	cmbs	*trbd*	*rmbs, corp*	t	-0.1923	6.7365
	rmbs	*vwst*	*corp, trbd*	t	0.0927	11.5838
4	*trbd3*	*rmbs*	*trbd6, fnbd, cmbs*	t	-0.0679	12.0842
	trbd6	*corp*	*fnbd, cmbs, rmbs*	I	-	-
	fnbd	*trbd*	*cmbs, rmbs, corp*	I	-	-
	cmbs	*vwst*	*rmbs, corp, trbd*	I	-	-
5	*trbd3*	*corp*	*trbd6, fnbd, cmbs, rmbs*	I	-	-
	trbd6	*trbd*	*fnbd, cmbs, rmbs, corp*	I	-	-
	fnbd	*vwst*	*cmbs, rmbs, corp, trbd*	I	-	-
6	*trbd3*	*trbd*	*trbd6, fnbd, cmbs, rmbs, corp*	t	0.0370	27.4600
	trbd6	*vwst*	*fnbd, cmbs, rmbs, corp, trbd*	t	-0.0743	11.9488
7	*trbd3*	*vwst*	*trbd6, fnbd, cmbs, rmbs, corp, trbd*	F	0.2482	-

The estimation results for the D-vine copula are presented in Table 5.18. The lack of a root variable in each tree is a key difference between this estimation and that done for the C-vine copula. Instead, the pairs are arranged in the fashion of a line where the first variable pairs with the second, the second goes on to pair with the third, and so on. Again, a popular copula for the pairs is the Student-*t* copula with progressively smaller degrees of dependence (and more degrees of freedom) as we move down to later trees with larger conditioning sets. In fact, quite a number of pairs in the fourth and fifth trees do not even reject the null hypothesis of independence and are modeled with an independence copula. One interesting observation is the *trbd3-cmbs* pair in the third tree. Conditional on six-month U.S. Treasury bills and non-U.S. government bonds, the model estimates negative dependence between these two fixed income asset classes. It may be due to very negative correlations during the financial crisis when U.S. Treasury securities, including bills, were a safe haven and commercial mortgage-backed securities were anything but that.

We proceed to simulate life insurance company asset returns that reflect the underlying dependencies as modeled by the C-vine and D-vine copula estimations in addition to the marginal distributions modeled by the GPD estimations. Regarding the number of returns to be simulated, we could produce any number of returns as long as it was computationally feasible. However, in order to match our original dataset, we will generate 4,234 returns for each asset class, which is the same number of historical returns in the common date range.

To build portfolios of these simulated returns, we use the same Combined Account weights as the risk hyperplane analysis. Recall that the weights for each reference year were determined based on a few elements. First, we hold fixed the actual Separate Account weights from that reference year. Second, we vary either the General Account equity or the General

Account corporate bond weights in order to focus the analysis on the company's allocation choice of one of these important asset classes. Third, we hold fixed relative proportions in either the non-equity or the non-corporate bond segment of the General Account. Again, these relative proportions are based on the actual industry-wide weights from the corresponding reference year. The equity weights are chosen from a range of 1% to 5% both to reflect the actual weights chosen by life insurers and the regulatory constraints on devoting a significant part of the General Account to high-risk securities like equities. As with the risk hyperplane analysis, the vine copula analysis appears to be quite partial to corporate bonds. Constraining the analysis to choose weights within a range more consistent with the actual weights chosen by the industry would produce certainly uninteresting results. The chosen weight would always be the maximum point in the range, and we would not the ability to see how different policyholder allocation decisions factored in to the company's allocation decisions. So, we use a range of 80% to 100% to choose an optimal corporate bond General Account weight.

For each portfolio formed, we calculate a number of statistics. First, we define a left tail by calculating an estimate of the Value-at-Risk by finding the portfolio return that corresponds to various threshold points including 1%, 2.5%, 5.85%, and 10%. We choose to estimate the VaR at 5.85% instead of the more common point of 5% because 248 observations, which is the number of observations used in the risk hyperplane analysis to define the portfolio's left tail, corresponds to approximately 5.85% of the full time series of portfolio returns. Finally, we estimate the Sortino Ratio for each portfolio-threshold combination by estimating the second lower partial moment for the returns in the left tail and the difference between the average portfolio return and the VaR. The Sortino Ratio estimates are presented in Tables 5.19 and 5.20, respectively, for equities and in Tables 5.21 and 5.22, respectively, for corporate bonds.

Table 5.19. C-Vine Sortino Ratios by Reference Year and Equity Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. The portfolio returns are determined by the Combined Account weights used and returns simulated for each asset class based on the joint dependencies modeled by a C-vine copula. The weights depend on the actual Separate Account, the non-equity General Account allocations for each reference year, and the General Account equity weight from the column heading. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The target rate is based on the Value-at-Risk calculated at various threshold points. The maximum Sortino Ratio for each reference year and VaR threshold is in the "Optimal" column, and the corresponding equity weight is in parentheses.

Panel A – 2013						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	1.5400	1.5370	1.5376	1.5389	1.5292	1.5419 (4.40%)
VaR Threshold of 2.50%	1.3139	1.3061	1.3040	1.3045	1.2954	1.3160 (1.40%)
VaR Threshold of 5.85%	1.1123	1.1000	1.0891	1.0904	1.0901	1.1123 (1.00%)
VaR Threshold of 10.00%	0.9075	0.9105	0.9069	0.9042	0.9009	0.9128 (1.60%)
Panel B – 2010						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	1.5031	1.5060	1.5122	1.5178	1.5269	1.5269 (5.00%)
VaR Threshold of 2.50%	1.3761	1.3579	1.3345	1.3252	1.3020	1.3761 (1.00%)
VaR Threshold of 5.85%	1.1345	1.1318	1.1321	1.1157	1.1065	1.1410 (2.40%)
VaR Threshold of 10.00%	0.9139	0.9097	0.9107	0.9133	0.9082	0.9144 (3.80%)
Panel C – 2006						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	1.5192	1.5208	1.5221	1.5298	1.5456	1.5456 (5.00%)
VaR Threshold of 2.50%	1.3220	1.3055	1.3076	1.3017	1.2973	1.3220 (1.00%)
VaR Threshold of 5.85%	1.1161	1.1127	1.1063	1.0958	1.0972	1.1163 (1.20%)
VaR Threshold of 10.00%	0.9076	0.9070	0.8990	0.8952	0.9000	0.9093 (2.20%)
Panel D – 2002						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	1.5237	1.4834	1.4496	1.4434	1.4334	1.5237 (1.00%)
VaR Threshold of 2.50%	1.4083	1.4215	1.4188	1.4163	1.4416	1.4416 (5.00%)
VaR Threshold of 5.85%	1.0850	1.0852	1.0755	1.0865	1.0834	1.0866 (1.40%)
VaR Threshold of 10.00%	0.9093	0.9033	0.9033	0.8951	0.8895	0.9093 (1.00%)
Panel E – 1998						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	1.5059	1.4947	1.4876	1.4832	1.4775	1.5059 (1.00%)
VaR Threshold of 2.50%	1.4173	1.3991	1.4010	1.3957	1.3784	1.4173 (1.00%)
VaR Threshold of 5.85%	1.1372	1.1367	1.1279	1.1284	1.1353	1.1418 (1.20%)
VaR Threshold of 10.00%	0.9316	0.9232	0.9270	0.9227	0.9201	0.9323 (3.40%)
Panel E – 1994						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	2.0401	1.9860	1.9813	1.9649	1.9525	2.0401 (1.00%)
VaR Threshold of 2.50%	1.7408	1.6960	1.7100	1.6554	1.5713	1.7408 (1.00%)
VaR Threshold of 5.85%	1.3535	1.3280	1.3068	1.2864	1.2793	1.3662 (1.40%)
VaR Threshold of 10.00%	1.0076	1.0082	1.0040	1.0025	1.0048	1.0200 (2.20%)

Table 5.20. D-Vine Sortino Ratios by Reference Year and Equity Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. The portfolio returns are determined by the Combined Account weights used and returns simulated for each asset class based on the joint dependencies modeled by a D-vine copula. The weights depend on the actual Separate Account, the non-equity General Account allocations for each reference year, and the General Account equity weight from the column heading. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The target rate is based on the Value-at-Risk calculated at various threshold points. The maximum Sortino Ratio for each reference year and VaR threshold is in the "Optimal" column, and the corresponding equity weight is in parentheses.

Panel A – 2013						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	2.1371	2.1135	2.0885	2.0906	2.0922	2.1371 (1.00%)
VaR Threshold of 2.50%	1.5468	1.5410	1.5326	1.5286	1.5282	1.5468 (1.00%)
VaR Threshold of 5.85%	1.2009	1.2007	1.2032	1.2057	1.2048	1.2063 (4.40%)
VaR Threshold of 10.00%	1.0229	1.0282	1.0289	1.0173	1.0188	1.0316 (2.60%)
Panel B – 2010						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	2.1886	2.1861	2.1591	2.1576	2.1826	2.1886 (1.00%)
VaR Threshold of 2.50%	1.5510	1.5393	1.5361	1.5272	1.5239	1.5510 (1.00%)
VaR Threshold of 5.85%	1.2480	1.2258	1.2096	1.2142	1.2142	1.2480 (1.00%)
VaR Threshold of 10.00%	1.0367	1.0309	1.0264	1.0253	1.0239	1.0367 (1.00%)
Panel C – 2006						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	2.1226	2.1108	2.1092	2.0870	2.0758	2.1226 (1.00%)
VaR Threshold of 2.50%	1.5548	1.5550	1.5415	1.5319	1.5476	1.5557 (1.80%)
VaR Threshold of 5.85%	1.2123	1.2033	1.1964	1.1966	1.1945	1.2123 (1.00%)
VaR Threshold of 10.00%	1.0305	1.0360	1.0357	1.0257	1.0252	1.0384 (2.40%)
Panel D – 2002						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	2.1362	2.1193	2.1433	2.1595	2.1668	2.1668 (5.00%)
VaR Threshold of 2.50%	2.0586	2.0668	2.0622	2.0115	1.9718	2.0682 (1.60%)
VaR Threshold of 5.85%	1.3656	1.3414	1.3262	1.3192	1.3169	1.3656 (1.00%)
VaR Threshold of 10.00%	1.0911	1.0707	1.0583	1.0404	1.0239	1.0911 (1.00%)
Panel E – 1998						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	2.0636	2.0717	2.1129	2.1496	2.1666	2.1697 (4.80%)
VaR Threshold of 2.50%	1.5466	1.5407	1.5240	1.5366	1.5255	1.5492 (1.20%)
VaR Threshold of 5.85%	1.2423	1.2392	1.2330	1.2309	1.2275	1.2423 (1.00%)
VaR Threshold of 10.00%	1.0180	1.0156	1.0199	1.0247	1.0248	1.0269 (4.40%)
Panel E – 1994						
Equity Weight	*1.00%*	*2.00%*	*3.00%*	*4.00%*	*5.00%*	*Optimal*
VaR Threshold of 1.00%	1.6545	1.6835	1.6381	1.6429	1.6025	1.7035 (1.60%)
VaR Threshold of 2.50%	1.5663	1.5350	1.5017	1.5294	1.5235	1.5669 (1.20%)
VaR Threshold of 5.85%	1.3880	1.3720	1.3465	1.3500	1.3409	1.3880 (1.00%)
VaR Threshold of 10.00%	1.1549	1.1488	1.1471	1.1316	1.1252	1.1549 (1.00%)

Table 5.21. C-Vine Sortino Ratios by Reference Year and Corporate Bond Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. The portfolio returns are determined by the Combined Account weights used and returns simulated for each asset class based on the joint dependencies modeled by a C-vine copula. The weights depend on the actual Separate Account, the non-corporate bond General Account allocations for each reference year, and the General Account corporate bond weight from the column heading. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The target rate is based on the Value-at-Risk calculated at various threshold points. The maximum Sortino Ratio for each reference year and VaR threshold is in the "Optimal" column, and the corresponding corporate bond weight is in parentheses.

Panel A – 2013						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.3115	1.3276	1.3371	1.3509	1.3769	1.3769 (100%)
VaR Threshold of 2.50%	1.2184	1.2164	1.2306	1.2164	1.2109	1.2348 (91%)
VaR Threshold of 5.85%	0.9853	0.9847	0.9848	0.9907	1.0112	1.0112 (100%)
VaR Threshold of 10.00%	0.8605	0.8532	0.8514	0.8467	0.8464	0.8605 (80%)
Panel B – 2010						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.3930	1.4180	1.4369	1.4395	1.4351	1.4421 (94%)
VaR Threshold of 2.50%	1.2116	1.1991	1.2030	1.2044	1.2050	1.2171 (81%)
VaR Threshold of 5.85%	0.9891	0.9950	0.9961	0.9975	1.0054	1.0054 (100%)
VaR Threshold of 10.00%	0.8534	0.8554	0.8575	0.8564	0.8548	0.8581 (92%)
Panel C – 2006						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.3347	1.3688	1.4020	1.4075	1.4115	1.4115 (100%)
VaR Threshold of 2.50%	1.2165	1.2148	1.2053	1.2089	1.2030	1.2216 (81%)
VaR Threshold of 5.85%	0.9891	0.9855	0.9964	1.0009	1.0099	1.0103 (99%)
VaR Threshold of 10.00%	0.8450	0.8498	0.8522	0.8520	0.8612	0.8629 (99%)
Panel D – 2002						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.4901	1.4756	1.4446	1.4748	1.5189	1.5189 (100%)
VaR Threshold of 2.50%	1.3938	1.3854	1.3918	1.3938	1.4005	1.4005 (100%)
VaR Threshold of 5.85%	1.1101	1.1296	1.1397	1.1430	1.1615	1.1615 (100%)
VaR Threshold of 10.00%	0.9190	0.9283	0.9398	0.9436	0.9564	0.9564 (100%)
Panel E – 1998						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.4648	1.4531	1.4405	1.3856	1.4013	1.4648 (80%)
VaR Threshold of 2.50%	1.2057	1.2112	1.2099	1.2294	1.2386	1.2415 (99%)
VaR Threshold of 5.85%	0.9873	0.9895	1.0103	1.0347	1.0261	1.0386 (96%)
VaR Threshold of 10.00%	0.8602	0.8428	0.8495	0.8492	0.8620	0.8620 (99%)
Panel E – 1994						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.6374	1.7289	1.8336	1.8855	1.8579	1.8927 (94%)
VaR Threshold of 2.50%	1.4238	1.4903	1.5394	1.5633	1.5944	1.5944 (100%)
VaR Threshold of 5.85%	1.1694	1.1856	1.2056	1.2626	1.2835	1.2893 (98%)
VaR Threshold of 10.00%	0.9106	0.9287	0.9658	0.9824	1.0261	1.0261 (100%)

Table 5.22. D-Vine Sortino Ratios by Reference Year and Corporate Bond Weight

The table contains the Sortino Ratio estimates for portfolios of the includible asset classes. The portfolio returns are determined by the Combined Account weights used and returns simulated for each asset class based on the joint dependencies modeled by a D-vine copula. The weights depend on the actual Separate Account, the non-corporate bond General Account allocations for each reference year, and the General Account corporate bond weight from the column heading. The Sortino Ratio is defined to be the ratio of the portfolio's return in excess of the target rate and tail risk as measured by the second lower partial moment. The target rate is based on the Value-at-Risk calculated at various threshold points. The maximum Sortino Ratio for each reference year and VaR threshold is in the "Optimal" column, and the corresponding corporate bond weight is in parentheses.

Panel A – 2013						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.9658	1.9750	1.9184	1.8300	1.9077	1.9826 (84%)
VaR Threshold of 2.50%	1.6838	1.6742	1.6943	1.7371	1.7235	1.7371 (95%)
VaR Threshold of 5.85%	1.3431	1.3456	1.3363	1.3601	1.3588	1.3639 (97%)
VaR Threshold of 10.00%	1.0799	1.0954	1.0997	1.0835	1.0876	1.1031 (89%)
Panel B – 2010						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.9290	1.9210	1.8739	1.9075	2.0232	2.0232 (100%)
VaR Threshold of 2.50%	1.7424	1.7566	1.7673	1.8007	1.8315	1.8315 (100%)
VaR Threshold of 5.85%	1.3614	1.3826	1.3887	1.4157	1.4249	1.4249 (100%)
VaR Threshold of 10.00%	1.1104	1.1152	1.1248	1.1535	1.1434	1.1535 (95%)
Panel C – 2006						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.9617	1.9356	1.8715	1.8900	1.9296	1.9617 (80%)
VaR Threshold of 2.50%	1.7327	1.7493	1.7549	1.7772	1.8032	1.8032 (100%)
VaR Threshold of 5.85%	1.3396	1.3549	1.3586	1.3692	1.4067	1.4067 (100%)
VaR Threshold of 10.00%	1.0952	1.0964	1.0994	1.1215	1.1139	1.1288 (98%)
Panel D – 2002						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	2.5267	2.5988	2.6271	2.7672	2.8257	2.8257 (100%)
VaR Threshold of 2.50%	1.9933	2.0139	2.0912	2.1832	2.2700	2.2700 (100%)
VaR Threshold of 5.85%	1.5651	1.5817	1.6516	1.6573	1.7005	1.7005 (100%)
VaR Threshold of 10.00%	1.1689	1.2047	1.2166	1.2232	1.2438	1.2438 (100%)
Panel E – 1998						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	1.8951	1.9314	1.9694	2.0952	2.2129	2.2129 (100%)
VaR Threshold of 2.50%	1.8179	1.8485	1.8845	1.8619	1.8726	1.8845 (90%)
VaR Threshold of 5.85%	1.4312	1.4660	1.4700	1.4567	1.4497	1.4701 (89%)
VaR Threshold of 10.00%	1.1349	1.1550	1.1673	1.1616	1.1477	1.1717 (92%)
Panel E – 1994						
Corporate Bond Weight	*80.00%*	*85.00%*	*90.00%*	*95.00%*	*100.00%*	*Optimal*
VaR Threshold of 1.00%	3.7649	3.5841	3.5528	3.5266	3.3855	3.8294 (81%)
VaR Threshold of 2.50%	2.4892	2.4320	2.3885	2.4121	2.4148	2.4892 (80%)
VaR Threshold of 5.85%	1.6750	1.6851	1.6659	1.6624	1.6780	1.6999 (83%)
VaR Threshold of 10.00%	1.2984	1.3030	1.3152	1.2846	1.2847	1.3152 (90%)

Figure 5.4. Optimal and Actual General Account Equity Allocations (Vine Copulas)
This graph plots the optimal General Account equity allocations as determined by the C-vine and D-vine copula portfolio models as compared with the historical industry-wide allocations. These allocations are based on historical portfolio weights for Separate Account assets and non-equity General Account assets from various years, which are shown on the horizontal axis. The vertical axis denotes the percentage of the General Account allocated to equities.

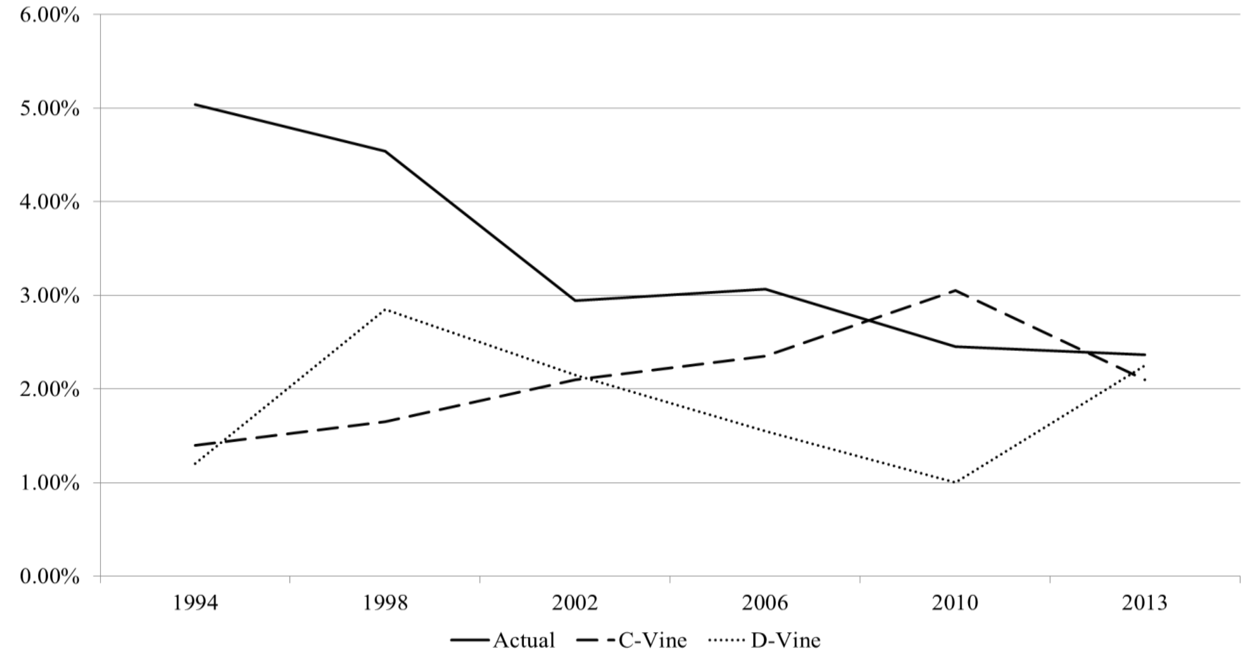

Reviewing the results for the equity allocation decision first, it is notable how much the vine copula models (both C-vine and D-vine versions) favor lower allocations to equities. In many of the reference year-threshold combinations, the optimal General Account allocation to equities is between 1% and 2%. In fact, the minimum allocation of 1% is chosen to be the optimal one with some regularity suggesting that perhaps we should expand the range of possible allocations below 1%. Thus, it is not surprising that there is a bit less variation from year to year in the average optimal allocations than with the risk hyperplane approach. This can be seen graphically in Figure 5.4, which shows the actual General Account equity allocations and the optimal C-vine and D-vine copula allocations averaged over the thresholds. In fact, the average optimal equity General Account allocations are almost entirely in the 1.0% to 2.5% range. The exceptions to this are when the C-vine model calls for an allocation of 3.05% in 2010 and when the D-vine model calls for an allocation of 2.85% in 1998.

We can also see from Figure 5.4 that the D-vine results deviate much further from the C-vine results than the GPD and GEVD models did in the risk hyperplane approach. This poses the question of whether the C-vine or D-vine copula model better models the underlying data. To answer this question, we conduct two tests that compare models, and they are the likelihood ratio test proposed by Vuong (1989) and the test proposed by Clarke (2007). Both of these are based on ratios of the competing vine copula densities and reject the null hypothesis of indistinguishability if the ratios are sufficiently large across our dataset. Both of them can also be corrected for the number of parameters used. Initially, we performed the Vuong test on the two models, and the C-vine model appeared to be the better model based on the direction of the test statistic, but the p-values were in the 15-19% range depending on which correction method used. Second, we performed the Clarke test, and this strongly suggested that the C-vine model is to be preferred with a p-value less than 0.10%. To confirm this, we looked at the log-likelihood values for the vine copula model estimations. The log-likelihood value for the C-vine model is 1,116.436 while it is -3,658.303 for the D-vine model, which is clearly smaller than that of the C-vine model. Hence, we will focus our analysis of the vine copula results on those produced by the C-vine model.

However, the biggest difference between the vine copula results for equities and those for the risk hyperplane approach is in the trend. Recall that the optimal risk hyperplane allocations, as shown in Figure 5.2, had an overall downward trend as we move from 1994 allocations to 2013 allocations. These optimal General Account equity allocations appear to have an upward trend. If we look closer at Figures 5.2 and 5.4, though, we will notice that the main difference in the results lies in the optimal allocations for 1994 and 1998. With the risk hyperplane approach, the optimal General Account equity allocations were in the 3.0% to 4.5% range while these are below 2%. Thus, the vine copula models must be modeling the joint dependencies in such a way that these earlier sets of allocations require lower allocations to equities. Reviewing the tail risk measurements of the C-vine model portfolios reveals that the square root of the lower partial moment is generally higher under the vine copula model than with the risk hyperplane analysis. The average tail risk for the vine copula model is about 0.32% and 0.53% in 1994 and 1998, respectively, while it is only about 0.16% and 0.30% for the same years in the GPD risk hyperplane model. So, this is a possible explanation for these relatively stark differences between the two approaches in the early years of our analysis.

Different tail risk and portfolio return sensitivities to changes in the General Account's equity allocation are another possible explanation. We can see this by reviewing the excess returns and tail risk measurements for the portfolios based on equity weights of 1% and 5%. With the risk hyperplane GPD model, the average tail risk based on the square root of the second lower partial moment is about 0.29% with an equity weight 1% and about 0.32% with a weight of 5%, which is an increase of about three basis points. The corresponding tail risk measurements for the C-vine copula model are 0.51% and 0.56%, which is an increase of about five basis points. So, the C-vine model appears to be more sensitive to increases in the equity weight in terms of tail risk.

For excess returns, the risk hyperplane approach appears to be more sensitive to increases in equity weight. In this model, the average excess return is about 0.80% with an equity weight of 1% and about 0.88% with a weight of 5% for an increase of about eight basis points. In the C-vine copula model, the corresponding excess return measurements are 0.65% and 0.70% for an increase of about five basis points. Given that a Sortino Ratio has excess returns in the numerator and tail risk in the denominator, it makes sense for the model that has faster increases in the tail risk and slower increases in the returns to prefer lower equity weights.

It is also interesting that the industry appears to have gradually caught up to optimal allocations suggested by the vine copula models. Perhaps life insurers should have been investing in equities closer to the current levels for most of the past twenty years. This "closing of the gap" effect actually occurs in both modeling approaches, but, as we just discussed, the main difference is the vine copula models start with a bigger gap. A final observation of Figure 5.4 that is worth mentioning is that both the C-vine and D-vine copula models also suggest that the industry's 2013 General Account equity allocation is either optimal or very close to it. This is one point on which both models seem to highly agree.

Lastly, we review the vine copula results of the corporate bond General Account allocation decision in Tables 5.21 and 5.22. The results from the vine copula models are largely in line with those from the risk hyperplane models. If anything, the optimal allocations from the C-vine copula model are, on average, slightly less than those from the risk hyperplane model, but there is a lot of overlap as well. Both approaches recommend investing nearly all of the General Account in corporate bonds, which is consistently the case as we change the exogenous allocations. To be sure, the time period from which we have sampled our returns is a beneficial one for corporate bonds. Overall, investment-grade corporate bonds, as an asset class, have earned a return only slightly less than equities but with much less risk, both in terms of standard deviation and tail risk. Unfortunately, we are constrained from using more historical returns data as long as we try to model as many asset classes as possible. Recall that some of our asset classes, especially *cmbs*, do not have as long of an available history. It would be interesting to extend this study by bringing in some other historical time periods in which corporate bonds did not out-perform equities as much on a risk-adjusted basis.

H. Portfolio Performance

To conclude our analysis, we briefly test some of the portfolio performance implications of the optimal allocations produced by our risk hyperplane and vine copula models. To do so, we compare how a portfolio based on our optimal weights performs relative to one based on the actual historical weights of the industry. Admittedly, this is merely an in-sample test of performance, which biases us in favor of finding that our optimal weights produce superior performance. However, it should still provide some insights into how detrimental are the industry's deviations from our optimal allocations.

We use all daily returns in our data sample for the time period common to all asset classes, which is January 1998 to October 2014. Portfolio weights are held constant for four years, with monthly re-balancing, to match the four year gaps between the allocations used in the models. Weights for a particular reference year are applied to the portfolios starting in the year following the reference year (e.g., the allocations from 1998 start being used by the portfolios at the beginning of 1999). For 1998, the first year of daily returns in our sample, the weights are those from 1994. All of the allocations used are on a Combined Account basis, and returns in excess of the daily risk-free rate (from Kenneth French's website) are used for performance measurement.

The portfolio performance results are presented in Table 5.23. A number of performance statistics are presented for each portfolio. The portfolios include one based on the actual industry-wide asset allocations ("Actual") and a number of competing portfolios based on the optimal allocations from our models. These include the GPD and GEVD risk hyperplane models and the C-vine and D-vine copulas models. For each of these, there is an equity-focused version and a corporate bond-focused version. The statistics include the average daily log excess return relative to the corresponding portfolio's left tail threshold, the square root of the second lower partial moment, and the Sortino Ratio. The left tail threshold is set equal to the 248th most negative daily return for each portfolio. This corresponds to the 5.85% VaR estimate used in the vine copula analysis and the threshold used in the GPD modeling for the risk hyperplane analysis. This same threshold is used in the calculation of the second lower partial moment, which is calculated based on the ex post version presented in Section V.C.1.

Table 5.23. Portfolio Performance

The table contains the in-sample performance of several portfolios of life insurance company investment assets. The benchmark portfolio is based on the actual asset allocations of the industry. Against this benchmark, we compare several portfolios based on the optimal allocations produced by each of the models in this study. These include the GPD and GEVD risk hyperplane models and the C-vine and D-vine copula models. With each of these, we also have a version focused on the General Account equity weight decision and a version focused on the General Account corporate bond ("CB") weight decision. A number of performance statistics are presented for each portfolio. These include the average daily log excess return relative to the corresponding portfolio's left tail threshold ("Excess Return"), the square root of the second lower partial moment ("Tail Risk"), and the Sortino Ratio.

	Actual	**GPD Equity**	**GEVD Equity**	**GPD CB**	**GEVD CB**	**C-Vine Equity**	**D-Vine Equity**	**C-Vine CB**	**D-Vine CB**
Excess Return (%)	0.5643	0.5554	0.5553	0.5540	0.5551	0.5474	0.5499	0.5537	0.5542
Tail Risk (%)	0.4940	0.4855	0.4859	0.4626	0.4598	0.4899	0.4798	0.4638	0.4624
Sortino Ratio	1.1423	1.1440	1.1428	1.1976	1.2074	1.1174	1.1460	1.1939	1.1986

The results in Table 5.23 make it clear how the optimal allocations improve the risk-adjusted performance of a life insurance company's Combined Account investment portfolio. They do so by reducing the portfolio's tail risk exposure since all of the portfolios based on our optimal allocations have average returns lower than that of the portfolio based on the actual weights. The reduction in the tail risk across our models averages about 4.11%, and this leads to an average increase of about 2.29% in the Sortino Ratio. Thus, our results here suggest that life insurers could potentially see some improvement in risk-adjusted portfolio performance with our optimal asset allocations, but the potential is still somewhat limited. It is also quite interesting to note that the portfolio based on the optimal equity allocations from the C-vine copula model actually underperforms relative to the "Actual" portfolio, even though it is an in-sample test. This leads us to conclude that the life insurers do not appear to deviate very significantly from optimal allocations as far as ex post performance is concerned.

Table 5.23 also provides some evidence for why our models exhibited a preference for corporate bonds. The portfolios that are heavily focused on corporate bonds in the General Account ("GPD CB," "GEVD CB," "C-Vine CB," and "D-Vine CB") have even lower tail risk without sacrificing much in return. Still, these corporate bond-focused portfolios did not perform as well in the middle of the recent financial crisis, which was a significant tail event. When we focused only on 2008, these four portfolios had negative risk premiums of about -20% for the year. Meanwhile, the "Actual" portfolio and the equity-focused optimal portfolios had negative risk premiums of about -18%. In fact, the equity-focused portfolios seemed to behave during the crisis as one might hope from a tail risk focused allocation decision. They had negative risk premiums that were all somewhat smaller than the -18.3% risk premium earned by the "Actual" portfolio. Even in a time period beneficial to corporate bonds, there appears to still be some benefits to diversification in a crisis.

VI. Conclusion

A. Summary

In this dissertation, we used extreme value statistics to study the measurement and modeling of market risk for a specific type of investment portfolio, which is that of a life insurance company. We chose this application to study our market risk analysis because life insurance companies are particularly prone to tail risk due to the guarantees they provide on their products, especially in the newer "variable" products that open up the company to greater downside equity exposure.

First, we reviewed the nature of the life insurance business including the historical development of the business, the types of risks faced by life insurance companies, the tools available to life insurance companies to manage these risks, and the regulatory environment in which life insurance companies operate. After performing this review, we used two modeling approaches to measure and model the market risk of a typical life insurance company's investment portfolio. This analysis is used to provide insights into optimal asset allocation choices for life insurers in light of their tail risk exposure. In this way, our models are related to the approach used by Roy (1952) rather than being based on the more customary mean-variance optimization approach.

The first approach, that of the risk hyperplane, is mechanically similar to a classical Markowitz (1952) approach but with the key difference that the risk-return tradeoff is made based on tail risk, rather than total risk. As a result, the second lower partial moment is used to measure the risk of a candidate portfolio and excess portfolio returns are determined relative to a target rate that will differ from the risk-free rate in general. The second approach makes use of the mathematical sophistication embedded in the emerging tool of vine copulas. Although vine copulas have a theoretical foundation stretching back to Sklar (1959), it is only in the past twenty years or so that the theory has developed to the point where it can be very useful in finance and insurance applications like the one focused on here. In this study, we use two vine copula structures, which are the so-called C-vine and D-vine copula structures, to model the joint dependencies present in a typical life insurance company's investment portfolio.

This analysis enables us to determine some key results. First, all of the modeling approaches (risk hyperplane, C-vine copula, and D-vine copula) agree that the industry's aggregate allocation to equities in the General Account as of 2013 is effectively optimal. Since this year marks the end of our analysis period, it remains to be seen whether or not the industry can maintain such optimality in their investment decisions. Second, there is some evidence, particularly in the risk hyperplane analysis, that the roughly 50% decline in General Account equity allocations is related to the significant increase in equity exposure from the Separate Account. Since the mid-1990s, not only has the equity allocation within the Separate Account increased but the Separate Account also comprises a greater proportion of the combined assets due to growth of "variable" products. This result highlights a complication of making optimal asset allocation decisions for an investor that has at least some of their risk exposure determined by the asset allocations of others. Interestingly, the vine copula models do not share this result, and it appears that it is due to detecting higher levels of tail risk and greater sensitivities to higher equity weights, especially in the early years of our analysis.

Finally, we studied the in-sample portfolio performance of our optimal allocations relative to the actual allocations chosen, in aggregate, by the industry. We find that the optimal allocations generally produce an increase in risk-adjusted performance of just over 2%, and this is primarily due to the superior performance of the corporate bond-focused portfolios.

B. Directions for Future Research

We do not seek to make the claim that this research is all-encompassing and understand that it is limited in nature. Certainly, there are areas in which this analysis could and ought to be extended and improved upon in the future. One key area for future improvement is the objective of the optimization done here. We seek to optimize the life insurance company's trade-off of portfolio return and tail risk, as measured by the Sortino Ratio. However, a more appropriate optimization for a life insurance company is really the capital surplus, which is the amount by the assets exceed the liabilities. This is also the company's buffer for unexpected losses. Maximizing the investment portfolio's return and minimizing the tail risk does not necessarily maximize the company's capital surplus. Our analysis is focused entirely on the asset side of the balance sheet. However, the liabilities of a life insurance company may also be sensitive to financial market movements. For one thing, the reserves that are required on the company's products will be based on present values of future expected payouts, which would generally change as prevailing interest rates fluctuate. For variable products where the policyholder's guarantee is often put-like in nature (i.e., the value of the guarantee and the company's liability increases as the policy's account value declines), a negative shock in the equity market could simultaneously reduce the company's assets and increase the liabilities, which would be a double-hit to the capital surplus. As a result, analyzing the capital surplus is a better, albeit much more difficult, method of assessing the full impact of market risk on the financial health of the life insurance company.

Another major extension would be to expand the number of key risks under analysis. Recall that the key risks faced by a life insurance company include insurance, market, credit, liquidity, operational, group, systemic, and regulatory risks. A more complete assessment of optimal asset allocation and tail risk in the context of a life insurance company would incorporate one or more of these other key risks. It may be especially useful to incorporate credit risk in a more complete analysis of the corporate bond allocation decision. It would be interesting to see if doing so makes the optimal allocations more consistent with actual industry-wide allocations.

It would also be helpful to conduct a more rigorous multivariate analysis of the asset allocation decision. This could potentially help us understand better the impacts of certain changes in policyholder behavior, such as decreasing allocations to equities during or after a market crash, affect the optimal allocation decisions of the company. Although we would expect to see some kind of counterbalancing effect here (i.e., the company should pull back on equities when policyholders invest more of their funds in equities and vice versa), we were not able to provide definitive evidence of this in the current study. In addition, a more rigorous analysis of the portfolio performance implications, including out-of-sample tests, would provide some more insight into the cost of deviating from optimality.

Other areas of improvement are more technical in nature. For example, we only considered the C-vine and D-vine copula structures of which the C-vine copula had a better fit to our set of asset classes. A generalization of the C-vine and D-vine copula structures is called a regular vine ("R-vine") copula. It allows for multiple nodes in the same tree from which various paths of variable pairs can extend. Thus, a C-vine copula is a special case by restricting to one the number of nodes in each tree and by requiring all paths to extend from that central node. A D-vine copula is a special case by requiring all variable pairs to lie along the same path instead of allowing certain elements in the path to act as the node of one or more side paths. Although more complicated, an R-vine copula structure is more flexible given this less restrictive structure, so it could potentially be a better model than the structures used in our study.

In addition, we were generally drawing conclusions from a limited number of observations due to having a relatively small number of historical years to base those observations on. We could increase the statistical power of these conclusions by trying to expand the number of observations, although it may require more simulations to generate observations due to having no asset allocation data prior to the mid-1990s.

Finally, one question left unexplored by our study relates to the relative merits of using the vine copula and risk hyperplane approaches. Certainly, both approaches are feasible given the vast computational power available now to almost anyone with the knowledge of how to use it. It is also obvious that the two approaches involve very different levels of theoretical rigor. The risk hyperplane approach is a "brute force" method that helps determine what we are interested in knowing, which is the relationship between asset allocation, portfolio return, and portfolio tail risk, but in a somewhat unsophisticated manner. Vine copulas, on the other hand, have much greater mathematical sophistication and more rigorous theoretical foundations. This enables the investor to potentially gain a more sophisticated understanding of the joint dependencies between the assets in a portfolio, which is a key advantage for using vine copulas. It allows the practitioner to gain a deeper knowledge of the impact on the portfolio's tail risk exposure by adjusting the allocations. With the risk hyperplane approach, these insights come from a kind of "trial and error" method by which we see how the end result is affected by altering allocations rather than trying to directly model the underlying dependencies driving the results. Still, the simplicity and ease of understanding the method gives the risk hyperplane a different kind of advantage. Perhaps the computational power available now will enable a practitioner to gain the same insights that could be found theoretically with vine copulas. This is one area of possible future research that will allow us to better understand if the mathematical sophistication of vine copulas is worth the cost or ultimately unnecessary.

REFERENCES

REFERENCES

Acharya, Viral V., 2011, "Systemic Risk and Macro-Prudential Regulation," Monograph for the Asian Development Bank.

Allen, David E., Michael McAleer, and Abhay K. Singh, 2014, "Risk Measurement and Risk Modelling Using Applications of Vine Copulas" (No. TI 14-054/III), Discussion Paper, Tinbergen Institute.

American Council of Life Insurers, 2010, *Life Insurers Fact Book 2010*.

American Council of Life Insurers, 2014, *Life Insurers Fact Book 2014*.

Bedford, Tim, and Roger M. Cooke, 2001, "Probability Density Decomposition for Conditionally Dependent Random Variables Modeled by Vines," *Annals of Mathematics and Artificial Intelligence* 32, 245-268.

Bedford, Tim, and Roger M. Cooke, 2002, "Vines – A New Graphical Model for Dependent Random Variables," *The Annals of Statistics* 30, 1031-1068.

Bernanke, Ben S., "Implementing a Macroprudential Approach to Supervision and Regulation," 47th Annual Conference on Bank Structure and Competition, Federal Reserve Bank of Chicago, Chicago, 5 May 2011.

Bernstein, Peter, 1998, *Against the Gods: The Remarkable Story of Risk* (John Wiley & Sons, Inc., New York, NY).

Board of Governors of the Federal Reserve System, 2014, *Federal Reserve Statistical Release Z.1: Financial Accounts of the United States*.

Booth, G. Geoffrey, and John Paul Broussard, 2015, "The Sortino Ratio and the Generalized Pareto Distribution: An Application to Asset Allocation," *Wiley Handbook in Financial Engineering and Econometrics: Extreme Value Theory and its Applications to Finance and Insurance* (in press).

Brechmann, Eike C., and Claudia Czado, 2013, "Risk Management with High-Dimensional Vine Copulas: An Analysis of the Euro Stoxx 50," *Statistics & Risk Modeling* 30, 307-342.

Brechmann, Eike C., and Ulf Schepsmeier, 2013, "Modeling Dependence with C- and D-Vine Copulas: The **R** Package CDVine," *Journal of Statistical Software* 52, 1-27.

Browne, Sid, 1995, "Optimal Investment Policies for a Firm with a Random Risk Process: Exponential Utility and Minimizing the Probability of Ruin," *Mathematics of Operations Research* 20, 937-958.

Business Wire, "MetLife to Ask Federal Court to Review SIFI Designation," 13 Jan. 2015 <http://www.businesswire.com/news/home/20150113005135/en/MetLife-Federal-Court-Review-SIFI-Designation#.VconoP7bKUk>.

Carmona, Rene, 2014, *Statistical Analysis of Financial Data in R* (Springer Science+Business Media, New York, NY).

Castillo, Enrique, and Ali S. Hadi, 1997, "Fitting the Generalized Pareto Distribution to Data," *Journal of the American Statistical Association* 92, 1609-1620.

Chiu, M. C., and D. Li, 2009, "Asset-Liability Management Under the Safety-First Principle," *Journal of Optimization Theory and Applications* 143, 455-478.

Clarke, Kevin A., 2007, "A Simple Distribution-Free Test for Nonnested Model Selection," *Political Analysis* 15, 347-363.

Clendenin, William, 1932, *Brief Outline History of Life Insurance* (American Conservation Company, Chicago, IL).

Consiglio, Andrea, Antonio Pecorella, and Stavros A. Zenios, 2009, "A Conditional Value-at-Risk Model for Insurance Products with a Guarantee," *International Journal of Risk Assessment and Management* 11, 122-137.

Consiglio, Andrea, Flavio Cocco, and Stavros A. Zenios, 2008, "Asset and Liability Modelling for Participating Policies with Guarantees," *European Journal of Operational Research* 186, 380-404.

Cont, Rama, 2001, "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues," *Quantitative Finance* 1, 223-236.

Czado, Claudia, 2010, "Pair-Copula Constructions of Multivariate Copulas," in P. Jaworski, F. Durante, W. Härdle, T. Rychlik, eds.: *Copula Theory and Its Applications* (Springer-Verlag, Berlin).

Czado, Claudia, Ulf Schepsmeier, and Aleksey Min, 2012, “Maximum Likelihood Estimation of Mixed C-Vines with Application to Exchange Rates,” *Statistical Modelling* 12, 229-255.

Deng, Ling, Chaoqun Ma, and Wenyu Yang, 2011, “Portfolio Optimization via Pair Copula-GARCH-EVT-CVaR Model,” *Systems Engineering Procedia* 2, 171-181.

Dickey, David, and Wayne A. Fuller, 1979, “Distribution of the Estimates for Autoregressive Time Series with a Unit Root,” *Journal of the American Statistical Association* 74, 427-431.

Ernst & Young, 2014, “Contributions of the Life Insurance Industry to the US Economy: Financial Protection of Families, Employment and Household Income, and Long-Term Capital Investments.”

Fama, Eugene F., 1965, “Portfolio Analysis in a Stable Paretian Market,” *Management Science* 11, 404-419.

Federal Reserve Bank of New York, 2006, “Systemic Risk and the Financial System.”

Ferguson, Niall, 2008, *The Ascent of Money: A Financial History of the World* (The Penguin Press, New York, NY).

Financial Stability Oversight Council, 2013, “September 19, 2013, Meeting Minutes.”

Financial Stability Oversight Council, 2014, “December 18, 2014, Meeting Minutes.”

Graunt, John, 1665, *Natural and Political Observations Mentioned in a following Index, and made upon the Bills of Mortality* (William Hall, Oxford).

Hamilton, James D., 1994, *Time Series Analysis* (Princeton University Press, Princeton, NJ).

Hansmann, Henry, 1985, “The Organization of Insurance Companies: Mutual versus Stock,” *Journal of Law, Economics, & Organization* 1, 125-153.

Hart, Orson H., 1965, “Life Insurance Companies and the Equity Capital Markets,” *The Journal of Finance* 20, 358-367.

Holy Bible, The New American Bible, 2005, School and Church Edition (DeVore & Sons, Inc., Wichita, KS).

Hong, Yongmiao, Jun Tu, and Guofu Zhou, 2007, "Asymmetries in Stock Returns: Statistical Tests and Economic Evaluation," *The Review of Financial Studies* 20, 1547-1581.

International Association of Insurance Supervisors, 2003, "Stress Testing by Insurers Guidance Paper."

International Association of Insurance Supervision, 2013, *Insurance Core Principles, Standards, Guidance and Assessment Methodology*.

Jeng, Vivian, Gene C. Lai, and Michael J. McNamara, 2007, "Efficiency and Demutualization: Evidence from the U.S. Life Insurance Industry in the 1980s and 1990s," *The Journal of Risk and Insurance* 74, 683-711.

Joe, Harry, 1996, "Families of m-Variate Distributions with Given Marginals and $m(m-1)/2$ Bivariate Dependence Parameters," in L. Rüschendorf, B. Schweizer, M.D. Taylor, eds.: *Distributions with Fixed Marginals and Related Topics*, 120-141 (Institute of Mathematical Statistics, Hayward).

Joe, Harry, 1997, *Multivariate Models and Dependence Concepts* (Chapman & Hall, London).

Junior, Leonidas Sandoval, and Italo De Paula Franca, 2012, "Correlation of Financial Markets in Times of Crisis," *Physica A: Statistical Mechanics and its Applications* 391, 187-208.

Kurowicka, Dorota, and Roger M. Cooke, 2006, *Uncertainty Analysis with High Dimensional Dependence Modelling* (John Wiley & Sons, Chichester).

Kwiatkowski, Denis, Peter C.B. Phillips, Peter Schmidt, and Yongcheol Shin, 1992, "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root," *Journal of Econometrics* 54, 159-178.

Lamm-Tennant, Joan, 1989, "Asset/Liability Management for the Life Insurer: Situation Analysis and Strategy Formulation," *Journal of Risk and Insurance* 56, 501-517.

Lamm-Tennant, Joan, and Laura T. Starks, 1993, “Stock versus Mutual Ownership Structures: The Risk Implications,” *The Journal of Business* 66, 29-46.

Levy, Haim, and Ran Duchin, 2010, “Markowitz’s Mean-Variance Rule and the Talmudic Diversification Recommendation,” in John B. Guerard, Jr., ed.: *Handbook of Portfolio Construction*, 97-123 (Springer, New York, NY).

Liu, Chi Sang, and Hailiang Yang, 2004, “Optimal Investment for an Insurer to Minimize its Probability of Ruin,” *North American Actuarial Journal* 8, 11-31.

Longin, Francois, 2005, “The Choice of the Distribution of Asset Returns: How Extreme Value Theory Can Help?” *Journal of Banking & Finance* 29, 1017-1035.

Loretan, Mico, and Peter C.B. Phillips, 1994, “Testing the Covariance Stationarity of Heavy Tailed Time Series: An Overview of Theory with Applications to Several Financial Datasets,” *Journal of Empirical Finance* 1, 211-248.

Mandelbrot, Benoit, 1963, “The Variation of Certain Speculative Prices,” *The Journal of Business* 36, 394-419.

Markowitz, Harry, 1952, “Portfolio Selection,” *The Journal of Finance* 7, 77-91.

National Association of Insurance Commissioners, 2007, “Life Insurance Buyer’s Guide.”

National Association of Insurance Commissioners, 2010, “The United States Insurance Financial Solvency Framework.”

National Association of Insurance Commissioners, 2013, *State of the Life Insurance Industry: Implications of Industry Trends*.

National Association of Insurance Commissioners, 2014, “Top 25 Life/Fraternal Market Share.”

National Organization of Life & Health Insurance Guaranty Associations, 2014, “The Nation’s Safety Net: The Life & Health Insurance Guaranty Association System.”

Nelsen, Roger B., 2006, *An Introduction to Copulas*, 2nd ed. (Springer-Verlag, Berlin).

O'Donnell, Terence, 1936, *History of Life Insurance* (American Conservation Company, Chicago, IL).

Office of the Comptroller of the Currency, 2008, "A Guide to the National Banking System."

Organisation for Economic Co-operation and Development, 1998, *Insurance Statistics Yearbook.*

Organisation for Economic Co-operation and Development, 2006, *Insurance Statistics Yearbook.*

Organisation for Economic Co-operation and Development, 2013, *Insurance Statistics Yearbook.*

Patton, Andrew J., 2004, "On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation," *Journal of Financial Econometrics* 2, 130 – 168.

Roy, Andrew D., 1952, "Safety First and the Holding of Assets," *Econometrica* 20, 431-449.

Shakespeare, William, 1598, *The Merchant of Venice*, Joseph Pearce, ed. (Ignatius Press, San Francisco, CA, 2009).

Sharpe, William F., and Lawrence G. Tint, 1990, "Liabilities – A New Approach," *Journal of Portfolio Management* 16, 5-10.

Sherris, Michael, 1992, "Portfolio Selection and Matching: A Synthesis," *Journal of the Institute of Actuaries* 119, 87-105.

Sklar, M., 1959, "Fonctions de Répartition à n Dimensions et Leurs Marges," *Publications de l'Institut de Statistique de L'Université de Paris* 8, 229-231.

Singh, Vijay P., 1998, *Entropy-Based Parameter Estimation in Hydrology* (Kluwer Academic Publishers, Dordrecht).

Tarullo, Daniel K., "Macroprudential Regulation," Yale Law School Conference on Challenges in Global Financial Services, Yale Law School, New Haven, CT, 20 Sep. 2013.

The Courant, "City Nicknames In Connecticut," 2 Sep. 2014 <http://www.courant.com/ community/hc-city-nicknames-in-connecticut-pg,0,6393142.photogallery>.

Viswanathan, Krupa S., and J. David Cummins, 2003, "Ownership Structure Changes in the Insurance Industry: An Analysis of Demutualization," *The Journal of Risk and Insurance* 70, 401-437.

Vuong, Quang H., 1989, "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica* 57, 307-333.

Wehrhahn, Rodolfo, 2008, "Introduction to Reinsurance," *Primer Series on Insurance* (The World Bank, Washington, D.C.).