Equilibrium Risk Pools in a Regulated Market With Costly Capital

Stephen J. Mildenhall

Created 2020-07-15 21:36:01.733338

Abstract

This paper investigates equilibrium risk pools in a market with risk-based solvency regulation and costly capital. It considers a market with two classes of risk, each having different aggregate volatility characteristics, such as personal auto and catastrophe exposed property. It identifies three possible equilibrium solutions: a single multiline pool, a multiline pool and a monoline pool, and two monoline pools. It determines conditions under which each of the three solutions occurs. The requirements are sensitive to the relative tail risk of the two classes and the capital standard. The model shows that the more volatile risk class bears a higher proportion of the capital cost. The results help explain various features seen in insurance markets, including the structure of the Florida homeowners market and the US medical malpractice market, and it can be applied more broadly to any regulated risk market.

JEL Codes: G22, G10

1 Introduction

Risk is better shared. The famous mutuality principle states that diversifiable risk should be removed by pooling. Remaining non-diversifiable risk is borne by each agent in proportion to their share of aggregate risk tolerance, Eeckhoudt, Gollier, and Schlesinger (2011). These results produce a world where risk-averse agents optimally pool all their risk.

In the real world, there is generally no way to pool risks in the manner envisioned by theory, a fact that Arrow (1996) uses to explain the smaller than expected role of insurance in a modern economy. Practical insurance involves limited liability insurance pools whose payouts are state-independent. The pools must attract capital in a competitive market where it has a non-zero opportunity cost. As a result, risk pooling is expensive and has a cost above actuarial estimates, even before transaction and frictional costs.

The benefits of pooling lead us to expect that a few large multiline companies should dominate the insurance market. Indeed, large multiline carriers write a substantial proportion of business. But looking more closely, we see many deviations from the pooling-is-best hypothesis. There is incomplete pooling between personal lines and commercial lines. Commercial, which is a more volatile business, having higher limits and a more severe pricing cycle, tends to be written by smaller companies. And over 2,600 property-casualty insurers operate in the US. Monoline companies write a disproportionate share of volatile business, such as Florida homeowners business, California earthquake, and medical malpractice. (Note that monoline companies write US mortgage guaranty for legal reasons.) What explains this structure?

Cost-based factors, economies or dis-economies of scale, and management span and control provide one explanation, Cummins et al. (2010). These are undoubtedly important. But, since this is a question about risk it should have a risk-based answer, and that is what this paper seeks to provide.

In recent work, Ibragimov, Jaffee, and Walden (2018) state that basic structure questions in a risk market with one-sided protection remain unanswered. Working in a perfect market with frictional costs of capital they show that monoline solutions are more likely when risks are asymmetric or correlated and that multiline pooling solutions are more likely for well-behaved and independent risks. This paper comes to qualitatively similar results but uses a entirely different market model.

We show that under certain realistic assumptions when two risk classes buy insurance from capital-regulated, limited liability insurance pools facing costly capital, there are three possible market outcomes. There can be one pooled insurer, two monoline insurers, or one pool insurer and one monoline. One pool is the usual solution. Two monoline pools occur when the regulator capital function is super-additive, generally when insurance risk is very thick-tailed, Ibragimov and Walden (2007). In the third case, the higher volatility class splits between the pool and the monoline, and the lower risk class appropriates all of the diversification benefits. The higher risk class pays its monoline rate. This result is likely counter to insurance affordability and availability goals for the high risk class. These effects are more pronounced for weaker capital regulation and cheaper insurance.

The market structure solutions emerge from the interaction between the investor's view of risk and the regulator's. Even when the investor and the regulator risk measures are both coherent, their combination, which determines market prices, can fail to be subadditive. This failure results in monoline solutions. Our results extend Ibragimov, Jaffee, and Walden (2010), which had the critical insight that there is essentially only one marginal cost capital allocation consistent with pricing. Their marginal cost view becomes our marginal regulatory capital, and we replace their perfect market pricing model with an imperfect model.

A primary goal for this work is to answer market structure questions in a framework consistent with the real insurance market. Our title flags the two realities with which we are most concerned: regulation and costly capital. We address our approach to each in the next two subsections, starting with regulation.

1.1 Regulated Market

Our model assumes that insurance is compulsory and is bought for the protection of others. As a result, the buyer has no interest in its quality, and so solvency regulation is necessary to ensure the insurance requirement functions effectively, Cummins (1988).

Aon Benfield (2015) estimate that around 60 percent of global property-casualty premium is compulsory or quasi-compulsory. Broadly interpreted, compulsory insurance includes auto liability, much property insurance providing collateral protection, contractual general liability and surety, and workers compensation for employers, amongst other lines. Flood insurance, but curiously not earthquake, is required to obtain a mortgage in the US on an at-risk property. We use compulsory as a catch-phrase to include any legal, contractual, or doing-business requirements.

Generally, compulsory insurance must be bought from a licensed insurer, operating under risk-based capital regulation. US conforming mortgages require property insurance from an insurer with an adequate rating from a nationally recognized statistical rating organization.

We argue that insureds have no interest in the quality of their insurance either because they are judgment proof or because residual claims are covered by a guaranty fund or both. As a result, insureds are pure price buyers provided the insurance satisfies their compulsory insurance requirement.

1.2 Costly Capital

The theory of multiline pricing typically assumes a perfect complete market, but with a frictional cost of holding capital in an insurance company, Phillips, Cummins, and Allen (1998), Myers and Read (2001), Sherris (2006), Ibragimov, Jaffee, and Walden (2010), and Cummins (2000). These assumptions imply that diversifiable insurance risk is not priced. The catastrophe (cat) bond market is inconsistent with these implications. A cat bond is deliberately structured to minimize frictional costs: it has no management, it is domiciled in a tax-free jurisdiction, and is very lightly regulated. And yet cat bonds are typically priced at two to ten times best-estimate losses (see www.artemis.bm for examples of cat bond pricing). Investors charge for bearing risk beyond the cost of holding capital in insurance entities.

Which risks do investors price? All theories price systematic, non-diversifiable risk. And while hurricane or earthquake are conceivably systematic, the idea becomes less credible for minor perils covered by cat bonds at substantial margins. On the other hand, it is hard to argue that investors price pure roulette lottery risk since the mechanisms to pool and manage it are so evident. But this is a false dichotomy. Insurance is not a roulette lottery. In the phrase of Anscombe and Aumann (1963) it is an uncertain horse lottery, governed by unknown and ambiguous subject probabilities rather than unique objective ones. There is extensive evidence that investors are ambiguity averse, Ellsberg (1961). Underwriters consider a one in one hundred year event risky primarily because they do not have one hundred years of experience. Its objective probability is uncertain. Data is more abundant for more frequent events, and uncertainty is lower. Zhang (2002) and Klibanoff, Marinacci, and Mukerji (2005) describe ambiguity relevant to insurance pricing. The latter paper has been applied in an insurance context by Robert and Therond (2014), Dietz and Walker (2017) and Jiang, Escobar-Anel, and Ren (2020). Note that it is hard to distinguish risk aversion from ambiguity aversion because risky thick-tailed distributions are also more ambiguous.

We assume that the insurance market is imperfect. Investors are ambiguity averse but not necessarily risk-averse. Market prices incorporate investor ambiguity aversion using a non-additive pricing functional. Wang (1996) and Wang, Young, and Panjer (1997) apply a non-additive functional using distorted probabilities to insurance pricing, leveraging diverse theoretical underpinnings including Huber (1981), Schmeidler (1986), Schmeidler (1989), Yaari (1987), and Denneberg (1994). Within this theory, distortion risk measures (DRM) occur repeatedly and in many guises. Kusuoka (2001) characterize DRMs as coherent, law invariant, and comonotonic additive functionals. We assume investors price using a DRM. DRMs are also known as spectral risk measures, Acerbi (2002), because of their weighted-VaR representation.

DRMs have many appealing properties. However, they are not additive and include transaction costs, via an implied bid-ask spread, Castagnoli, Maccheroni, and Marinacci (2004). Important results from Chateauneuf, Kast, and Lapied (1996), De Waegenaere (2000), and especially Castagnoli, Maccheroni, and Marinacci (2002) and De Waegenaere, Kast, and Lapied (2003), show DRMs are consistent with general equilibrium models, even though they are non-additive. Perfect market pricing functionals are linear and hence additive, which means they do not allow a diversification benefit. DRMs are consistent with arbitrage-free pricing: the presence of transaction costs neutralizes apparent arbitrage opportunities caused by non-additivity.

A DRM is a special case of a Choquet integral. It is defined by a distortion function, which is an increasing concave function from $[0,1]$ to itself. Distortion functions price a binary insurance policy as a function of its probability of loss.

DRMs are practical and easy to work with, and, most important, they have a natural allocation of a pool's diversification benefit back to its members, which we outline in Section 3. The natural allocation relies on the fact that DRMs are coherent, law invariant and comonotonic additive.

DRMs respect diversification because they are coherent and hence subadditive. As a result, pooling two risks with no default will always be advantageous, provided both premiums are finite. Therefore the failure of pooling must rely on something more than a DRM pricing functional. The extra ingredient is the interaction between the pricing functional and solvency regulation.

Finally, we assume there are are no frictional costs for holding capital in an insurance entity. Cummins (2000) explains that agency conflict, tax, and regulation are the primary causes of frictional costs. Our model is structured to remove these costs: there are no taxes, there is no agency conflict because insurers are an autonomous contract with no discretion (like a cat bond), and regulation is limited to solvency requirements. There is no economic reason why investors have to transfer assets into the insurer entity at all; it could operate like an old-style Lloyds syndicate provided the regulator had sufficiently vicious rights to enforce payment.

1.3 Other Literature and Context

An extensive literature considers optimal risk-sharing between two risk classes. In contrast, we assume that insurance transfers risk to a separate investor group, independent of the risk owners. Ibragimov, Jaffee, and Walden (2018) call this one-sided protection. Broadly, the global insurance market operates in this manner. For example, Canadian pension funds are significant investors in US catastrophe insurance-linked securities: investors form a distinct group from risk owners. We ask whether it is most efficient for all insureds to combine into one multiline pool when accessing investors.

We can consider risk pooling as a reinsurance question. Doherty and Tinic (1981) suggest that reinsurance can increase the value of an insurer by decreasing its probability of default, thereby increasing the premium rates it can charge. Our model is consistent with this view, explicitly varying premium by default probability.

Cummins, Dionne, and Gagne (2008) shows that reinsurance is expensive but significantly reduces the volatility of loss ratio, reducing insolvency risk, and improving financial strength. Insurers accept these higher costs to reduce their underwriting risk. Insurers and pool management are focused on risk, and they drive the process of minimizing capital costs, whereas insureds only focus on price. In our model, insurance pools are a transparent pass through, but they have a legitimate interest in the shape of risk because it drives their cost of capital. Therefore, there is a potential role for reinsurance.

In our model, reinsurance pools premium rates rather than risk. If it were optimal to pool the risks, then they would be pooled directly in the first place. Differential premium rates can be an impediment to this pooling that reinsurance can help overcome. Reinsurance does not lead to a new equilibrium solution because the possibility of direct contracting makes reinsurance pools are unstable and, absent some enforcement mechanism, they unravel. Surprisingly pools with reinsurance have less risk pooling than those without it.

Our work is related to cost allocation. Tsanakas and Barnett (2003) computes an Aumann-Shapley cost allocation for a DRM. These costs provide cost signals to management to use to optimize their risk pools, assuming marginal revenue rates are constant. Ibragimov, Jaffee, and Walden (2010) had the insight that marginal cost to the pool drive allocations, whereas value to the insured drives marginal revenue. The two rarely align when a regulator-mandated risk measure controls capital. As a result, management will see an opportunity to modify pool composition to realize excess profits. This behavior is not allowed in our model because we assume perfect information, and insureds will not pay above fair market value for any insurance cash flow.

Our model is not concerned with individual risk tolerance; insureds purchased because insurance is compulsory. Insureds may or may not be risk-averse. For the same reason, we are not worried about the structure of optimal insurance policies.

1.4 Contribution and Contents

This paper offers the following contributions.

We show that it is practical to work in a realistic, imperfect market and obtain explicit premium, loss, and capital allocation results uniquely by using a DRM pricing measure.

We contribute to the problem of the equilibrium industry structure. We show that a market with two classes of risk will consist of one multiline carrier, two monolines, or a multiline and a monoline. This last split result generally obtains when one class is more risky than the other. In that case, the less risky class appropriates all the diversification benefits. These results depend on the capital standard and highlight that regulation can have unintended consequences.

We provide new insight into how limited liability and diversification interact. Limited liability produces subsidies between lines, which offset the benefits of diversification. These effects are more pronounced when the capital standard is weaker. In a dynamic model, they can introduce subsidies that disadvantage low-risk insureds. Regulation should consider these impacts when calibrating a capital standard.

Our methods can help insurers answer strategic questions related to geographic or class expansion. Regulation introduces a rigidity in rates that can discourage diversification and lock-in subsidies between classes, and our analysis makes it clear where and how such rigidities will manifest themselves.

Finally, our model provides a novel way of looking at relative volatility or tail risk by class. It shows that pooling behavior is very sensitively balanced, again emphasizing that a poorly chosen regulatory standard can have an outsized impact on the market.

The rest of the paper is structured as follows. Section 2 describes the market participants and their interactions. Section 3 recalls results giving the fair market value of insurance cash flows for each class in a multiple-class insurance pool. Section 4 derives the market outcomes that can occur. Section 5 gives some examples, and Section 6 concludes and suggests further research.

1.5 Notation and Conventions

The terminology describing risk measures is standard, and follows Föllmer and Schied (2011). We work on a standard probability space, Svindland (2009).

Total insured loss, or total risk, is described by a random variable X. X reflects policy limits but is not limited by insurer assets. $X = \sum_i X_i$ describes the split of losses by class. F, S, f, and q are the distribution, survival, density, and quantile function of X. $X \wedge a$ denotes $\min(X, a)$ and $X^+ = \max(X, 0)$. 1_A is the indicator function on a set A.

S, P, M and Q refer to expected loss, premium, margin and equity as a function of insurer assets, a. Premium equals loss plus margin; assets equal premium plus equity.

We use the actuarial sign convention: losses are positive and large positive values are bad.

2 Market Participants and Market Assumption

We work in a market with four participants: insureds, a regulator, insurance pools and investors, as shown in Figure 1.

We assume a one period model, with no expenses, no taxes, and a zero risk-free rate of interest. These are standard simplifying assumptions, e.g. Ibragimov, Jaffee, and Walden (2010).

At time $t = 0$ insureds form into limited liability insurance pools. The policies issued to pool participants aggregate to a total exposure $X = \sum_i X_i$. The regulator capital measure determines the amount of assets the pool needs to hold, $a = a(X)$. The pool raises a from a combination of premium and by selling its $t = 1$ residual value to investors to raise equity. At $t = 1$ claims become known and are paid. If losses $X \le a$ all insureds are paid in full and the residual value $a - X$ is distributed to investors. If $X > a$ the pool defaults and pays insureds on an equal priority, pro rata basis. The investors receive nothing. In both cases the pool is wound-up at $t = 1$. There are no transactions between 0 and 1 and hence no distinction between capital and assets.

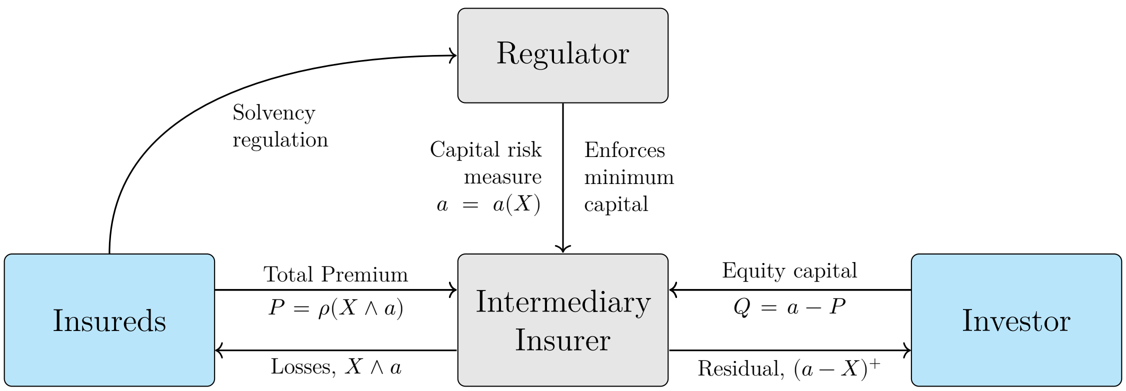

Figure 1: The four market actors and their interactions.

The investor prices with a DRM ρ corresponding to a distortion function g. The insurer purchases a quota share, with limit a, from the investors for a premium

$$\rho^a(X) := \rho(X \wedge a(X)). \tag{1}$$

Thus the market pricing functional combines investor and regulator risk functionals in an intricate and subtle way. All the effects we discuss are related to the properties of ρ^a. It is easy to see that if both a and ρ are positive homogeneous, monotone, translation invariant, law invariant, and comonotonic additive then ρ^a will be too. But, critically, ρ^a can fail to be subadditive even when both a and ρ are subadditive. It is this fact that makes the problem interesting and tricky.

Castagnoli, Maccheroni, and Marinacci (2004) shows DRMs always have a bid-ask spread. We do not count the spread as an expense because insurance positions are always long. Insurable interest laws make it impossible to short insurance.

We now specify the assumptions and behavior of each actor.

2.1 Insureds

There are two classes of insureds. Throughout, class 0 is a lower risk class and class 1 higher risk. Risk is relative. Insureds within a class can be considered identical, with the same insurance contract, the particulars of which are also irrelevant. As a result, adverse selection is not a problem.

We are only concerned with class risk and not with individual insured characteristics within the class. Correlations and dependencies between insureds determine class risk. It cannot be discerned insured by insured. Florida homeowners insurance is a high volatility class. Low limit, high frequency, and low severity US-style personal automobile is a low volatility class. An *individual* Florida homeowners policy may have the same range of outcomes as a home in France or Germany or Illinois. Still, the class risk of a portfolio of Florida homes will be higher because of the possibility of hurricane events.

Aggregate losses are homogeneous within each class, meaning that losses from a proportion x_i of class i has distribution $x_i X_i$ for fixed distributions X_i, $i = 1, 2$. A homogeneous loss model is a common approximation borrowed from finance, where x_i is a position size and X_i is a security price. It is used by Myers and Read (2001). X_i can be considered as the loss ratio for the class. Boonen, Tsanakas, and Wuthrich (2016) compares the homogeneous assumption to a more realistic compound Poisson model, and Mildenhall (2017) shows that a homogeneous assumption is not unrealistic for larger portfolios.

All market participants agree best-estimate subjective probabilities that underlying X_i. They are recognized as ambiguous and uncertain. Commercial catastrophe models provide such a pricing benchmark for the property catastrophe market.

Insureds purchase insurance because they face a compulsory insurance requirement. The requirement exists for the protection of third parties. Insureds are immune to the failure of their insurance, either because they are judgment-proof or because of a guaranty fund mechanism. As a result, insureds are pure price buyers of any insurance that fulfills their buying requirements. They are blind to the quality of the insurance they buy.

Individual risk aversion is not assumed in the model and is not necessary for the results.

2.2 Regulation

The solvency of insurance pools must regulated for compulsory insurance to be effective and provide the desired third-party protection.

Regulation takes the form of a regulatory capital risk measure a. We assume that a is law invariant, positive homogeneous, monotone, and translation invariant. In most cases a is given by value at risk (VaR) or tail value at risk (TVaR). US NAIC RBC and Solvency II are VaR-based; the Swiss Solvency Test is TVaR-based.

The regulatory risk measure is not assumed to be subadditive, and therefore may not be coherent.

There is no other regulation in the market. Rate regulation is not needed because we assume actors are well-informed. And there is no restriction on insurance pools, other than that they meet the capital requirement. A pool can consist of a single risk, or one pool could contain all risks, or any subset in-between. This flexibility means insureds always have the fall-back of being written in a pool of one, which caps their premium.

2.3 Insurer Intermediaries

Insurer intermediaries are transparent pass-through entities that provide compulsory insurance meeting the regulator's capital standard. They manage a pool of insureds and arrange for the risk to be transferred to investors. They are like a "smart-contract". They have no employees or management, and so there are no principle-agent problems.

A *pool* refers to a group of insureds. There are two types of pool. A monoline pool contains insureds from one class. A multiline pool contains insureds from more than one class.

An *insurance* pool is a limited liability insurance company. In theory, pools could contact directly with insurers and obtain unlimited cover. They will generally find this cost prohibitive in our model, but it is an option.

Although insurers are incorporeal, they have a real economic impact because they enable limited liability, which changes the cash flows paid to insureds. Limited liability is vital to ambiguity averse investors: it truncates the tail, and it lowers the ambiguity of the losses they assume.

We could argue that since investors ultimately bear all the risk they should be indifferent to how it is packaged into pools. This is not correct. The pools alter the actual payments made by investors because they impact the incidence of default. A multiline insurer structure will default in different states than two monoline insurers.

Limited liability requires a legal framework to function; this is why the market uses insurer legal entities rather than third-party pool managers who only act as brokers between insureds and investors but assume no risk. The limited liability contract binds the participants to equal priority in default. Marshall (2018) provides an illuminating discussion of the importance of a legal entity to implement limited liability in the context of the California Earthquake Authority.

The pool operates costlessly. Its only expense is the spread between subjective expected losses and the investor pricing functional. Forming pools to minimize this spread is the topic of the paper.

The insurance market is competitive. Since actors agree on probabilities, insureds only pay the fair price for insurance, and neither pools nor insurers make a profit.

Providing a genuine motivation for the existence of insurance pools is a significant advantage of our framework. Under perfect markets, there is no diversification benefit because the pricing rule is linear, and so there is no need for a pool. Justifying insurers usually involves an appeal to market access or transaction expenses, Ibragimov, Jaffee, and Walden (2010) (especially footnote 9 and surrounding discussion).

2.4 Investor Risk Bearers

Investors bear insurance risk. Risk is transferred from the insured group to investors, rather than being pooling between insureds.

A competitive market produces prices that are modeled by a DRM. Kusuoka (2001) and Föllmer and Schied (2011, chap. 4) show that a DRM is entirely characterized by a distortion function. A distortion function is an increasing concave function $g : [0,1] \to [0,1]$ satisfying $g(0) = 0$ and $g(1) = 1$. The DRM ρ_g associated with a distortion g acts on a non-negative random variable X as

$$\rho_g(X) := \int_0^\infty g(S(x))dx. \tag{2}$$

When ρ is combined with a regulatory capital function using eq. (6) we get the market pricing functional

$$\rho_g^a(X) = \rho_g(X \wedge a(X)) = \int_0^\infty g(S_{X\wedge a(X)}(x))dx = \int_0^{a(X)} g(S_X(x))dx \tag{3}$$

since the survival functions S_X and $S_{X\wedge a}$ agree on $[0, a)$, $S_{X\wedge a}(x) = 0$ for $x > a$ and $g(0) = 0$. We drop the subscript g from the notation below since g is fixed. Thus $\rho_g^a(X)$ is the premium charged for the insurance pool with total risk X.

2.5 Market Assumption Summary

Here is a summary of our insurance market assumptions.

1. Insureds are required to purchase compulsory insurance. They are pure price buyers.
2. Insurance losses are homogeneous within each pool.
3. Insurance must be purchased from a one period, limited liability insurance pool. Pools can consist of one or more policies from either or both classes. There is no restriction on the size or composition of pools. Single policy pools are allowed.
4. Insurance pools are required by regulation to capitalize according to a risk-based capital formula. It is law invariant, positive homogeneous, monotonic, translation invariant, but not necessarily subadditive. VaR is the archetype.
5. There is equal priority if the pool defaults because it has insufficient assets to pay its obligations.
6. Investors provide equity to insurance pools by buying its residual value. Investors price using a DRM ρ associated with a distortion g. Equation (3) computes the pool premium.
7. Insurance pools are costless to form and operate, other than the cost of risk transfer to investors. There are no taxes.
8. All market participants agree on a set of subjective probabilities. Insurance pools are transparent, and individuals know the fair market value of their insured claims. The market is perfectly competitive and acts to remove any excess profits.

Under these assumptions, we now investigate what types of equilibrium insurance pools will form. To do that, we start by deriving the fair market value of pool insurance to pool participants.

3 Market Pricing

This section describes what we call the natural allocation of pool premium, under the assumptions set out in section 2 The natural allocation has a long history. Delbaen (2000) shows it is a fair allocation in the sense of fuzzy games and has a directional derivative, marginal interpretation when ρ is differentiable. Equivalent formulations are derived several other papers, including Venter, Major, and Kreps (2006) and Tsanakas and Barnett (2003). It is consistent with Campi, Jouini, and Porte (2013), who show the rational price of X in a market with frictions must be anti-comonotonic with the state prices implied by X. In our application the signs are reversed: $g'(S(X))$ and X are comonotonic.

The natural allocation transfers perfect market pricing theory to a non-additive, imperfect market setting. The perfect market results are derived in Sherris (2006) and Ibragimov, Jaffee, and Walden (2010).

A perfect market pricing functional has the form $\mathsf{E}_{\mathsf{Q}}[\cdot]$, where Q is a risk-adjusted measure. The Radon Nikodym derivative of Q defines the state price density. Given losses by class $X = \sum_i X_i$ supported by assets a, the problem is to allocate the total premium $\mathsf{E}_{\mathsf{Q}}(X \wedge a)$ between classes. Specifying the payments each class receives in default is the critical step.

The most common approach to payments in default is called equal priority. It shares assets in proportion to unlimited claims. Individual insurance contracts promise to pay X_i to class i. In the aggregate, promised payments are split into realized payments and insurer default as

$$X = X \wedge a + (X - a)^+.$$

Multiplying both sides by X_i/X shows

$$\begin{aligned} X_i &= X_i \frac{X \wedge a}{X} + X_i \frac{(X-a)^+}{X} \\ &= (\text{payments to class } i) + (\text{default to class } i). \end{aligned}$$

The payments to class i under **equal priority** are defined as

$$X_i(a) := X_i \frac{X \wedge a}{X} = \begin{cases} X_i & X \le a \\ a\dfrac{X_i}{X} & X > a. \end{cases} \tag{4}$$

Under equal priority class i is paid in full when total losses are less than or equal to assets and payments are pro-rated down by X_i/X in default states. These payments are contractual and are legally specified. Note that $\sum_i X_i(a) = X \wedge a$.

Contractually specified payments to i are a function of X_i alone. These are the amounts promised to i by their insurance contract. $X_i(a)$ is the amount actually paid to i; it is a function of losses for all classes, $X_1, \dots, X_n$ as well as the total assets a held by the insurer. This distinction between promised and delivered payments is critical. It is responsible for almost all the complexity of insurance pricing. In economics, the benefit of a good to the buyer is usually independent of who else purchases it. With insurance risk pools that is not the case: the value to one insured is almost always changed in some way if the pool takes on other risks. To completely specify the payments to i it is necessary to know about possible payments on all other contracts. These facts have profound implications for pooling, as we shall see in the next section.

Sherris (2006) and Ibragimov, Jaffee, and Walden (2010) show that the fair premium for risk i is simply $\mathsf{E}_{\mathsf{Q}}[X_i(a)]$. It is an allocation of total premium $\mathsf{E}_{\mathsf{Q}}[X \wedge a]$ because expectations are linear.

Under a DRM, the state price density Q becomes a function of the risk being priced; it is the measure solving $\rho(X) = \mathsf{E}_{\mathsf{Q}}[X]$, subject to certain other conditions. But we can still use the perfect market approach to allocate premium: the fair premium for i is $\mathsf{E}_{\mathsf{Q}}[X_i(a)]$, even though Q is no longer fixed. We call this the **natural allocation premium**. There are some subtle points regarding whether Q is unique, but for DRMs they can be circumvented. The details of our approach are given in Major and Mildenhall (2020). Here we give a summary, describing explicit formulas to compute the natural allocation. Our market structure results follow from the insights that these formulas provide. The formulas rely on three auxiliary functions by class: α_i, β_i and κ_i.

The limited expected loss cost and premium for an insurance pool supported by assets a are given by

$$\mathsf{E}[X \wedge a] = \int_0^a S(x)\,dx, \tag{5}$$

$$\rho(X \wedge a) = \int_0^\infty g(S_{X\wedge a}(x))\,dx = \int_0^a g(S_X)(x)dx. \tag{6}$$

Using eq. (5) and the fundamental theorem of calculus we can interpret $S(a)$ as the loss density in the layer at a, that is, the derivative of $\mathsf{E}[X \wedge a]$ with respect to assets a. Similarly, by eq. (6), the premium density is $g(S(a))$. Premium and loss densities are the proportions of the limit of a thin layer attaching at a that are attributable to premium or loss. A **thin layer** at a means an insurance contract paying $1_{X>a}$, i.e., 1 when $X > a$ and 0 otherwise.

Expected losses for class i can be computed by conditioning on X as

$$\mathsf{E}[X_i(a)] = \mathsf{E}[\mathsf{E}[X_i(a) \mid X]] = \mathsf{E}[X_i \mid X \le a]F(a) + a\mathsf{E}[X_i/X \mid X > a]S(a). \tag{7}$$

Because of its importance in allocating losses, define

$$\alpha_i(a) := \mathsf{E}[X_i/X \mid X > a]. \tag{8}$$

The value $\alpha_i(x)$ is the expected proportion of recoveries by class i in a thin layer attaching at x. Since total assets available to pay losses always 100 percent of the layer and the chance the layer attaches is $S(x)$, it is intuitively clear $\alpha_i(x)S(x)$ is the derivative of $\mathsf{E}[X_i(a)]$ with respect to x. This is, in fact, the case by Major and Mildenhall (2020), Proposition 1. Thus

$$\mathsf{E}[X_i(a)] = \int_0^a \alpha_i(x)S(x)dx \tag{9}$$

and the class i loss density at x is $\alpha_i(x)S(x)$. Equation (9) gives a direct analog to eq. (5) for class i losses. It is remarkable, because it converts $X = \sum_i X_i$, a convolution of random variables, into a simple sum.

Total premium under ρ is given by eq. (6). We now compute the natural allocation premium and premium density for each class. Using integration by parts the price of an unlimited cover on X is

$$\rho(X) = \int_0^\infty g(S(x))\,dx = \int_0^\infty x g'(S(x)) f(x)\,dx = \mathsf{E}[Xg'(S(X)))]. \tag{10}$$

It is important that this integral is over all $x \geq 0$ so the $xg(S(x))|_0^a$ term disappears. This shows that $g'(S(X))$ is the Radon Nikodym derivative of Q discussed above. Thus the natural allocation premium is $\mathsf{E}[X_i g'(S(X))]$. It is denoted $\rho_X(X_i)$ to emphasize its dependence on X. The choice $g'(S(X))$ is economically meaningful because it weights the largest outcomes of X the most, which is appropriate from a social, regulatory and investor perspective. Since DRMs are comonotonic additive the same density can be used for $\rho(X \wedge a)$. This will not be true for non-comonotonic additive risk measures, Shapiro (2012).

Next, we define the premium risk-adjusted analog of the α_i as

$$\beta_i(a) := \mathsf{E}_{\mathsf{Q}}[X_i/X \mid X > a]. \tag{11}$$

$\beta_i(x)$ is the value of the recoveries paid to class i by a thin layer at a. By the properties of conditional expectations, we have

$$\beta_i(a) = \frac{\mathsf{E}[(X_i/X)Z \mid X > a]}{\mathsf{E}[Z \mid X > a]}. \tag{12}$$

The denominator equals $\mathsf{Q}(X > a)/\mathsf{P}(X > a)$. (Remember that while $\mathsf{E}_{\mathsf{Q}}[X] = \mathsf{E}[XZ]$, for conditional expectations $\mathsf{E}_{\mathsf{Q}}[X \mid \mathcal{F}] = \mathsf{E}[XZ \mid \mathcal{F}]/\mathsf{E}[Z \mid \mathcal{F}]$, see Föllmer and Schied (2011), Proposition A.12.)

To compute α_i and β_i we use the third function

$$\kappa_i(x) := \mathsf{E}[X_i \mid X = x], \tag{13}$$

the conditional expectation of loss by class, given the total loss. It is an important fact that the risk adjusted version of κ is unchanged because DRMs are law invariant.

Major and Mildenhall (2020) Theorem 2 now gives the following, essentially unique expression for the natural allocation of eq. (6)

$$\begin{aligned}\rho_{X\wedge a}(X_i(a)) = \mathsf{E}_{\mathsf{Q}}[X_i g'(S(X)) \mid X \leq a](1 - g(S(a)))+ \\ a\mathsf{E}_{\mathsf{Q}}[(X_i/X)g'(S(X)) \mid X > a]g(S(a)).\end{aligned} \tag{14}$$

Moreover, we get a premium analog of eq. (9)

$$\rho_{X\wedge a}(X_i(a)) = \int_0^a \beta_i(x) g(S(x))\,dx. \tag{15}$$

To recap: the premium formulas eqs. (14) and (15) have been derived assuming only that capital is provided at a cost g and there is equal priority by class. Both formulas are computationally tractable. There is no need to assume the X_i are independent. The formulas produce an entirely general, canonical determination of premium in the presence of shared costly capital. They extend Grundl and Schmeiser (2007), who pointed out that with an additive pricing functional there is no need to allocate capital to price, to the situation of a non-additive DRM pricing functional.

4 Market Outcomes

We now apply the general theory developed in section 3 to the specific questions introduced in section 2. First we look at some relevant implications of the formalism developed so far.

4.1 Properties of the Natural Allocation

The results in this subsection only assume we are pricing with a DRM and have equal priority in default. They do not assume a homogeneous loss model. Nor do they assume the risk classes are independent.

The **margin** in premium is defined as the difference between premium and expected loss. Margin is the cost of bearing risk. It is a measure of the spread of investor probabilities over best-estimate subjective probabilities. Equations (5), (6), (9) and (15) give the total margin and its natural allocation by class

$$
\begin{aligned}
M(a) &:= \int_0^a g(S(t)) - S(t)\,dt, \\
M_i(a) &:= \int_0^a \beta_i(t)g(S(t)) - \alpha_i(t)S(t)\,dt. \qquad (16)
\end{aligned}
$$

Differentiating recovers the corresponding margin densities

$$
M'(a) = g(S(a)) - S(a); \quad M_i'(a) = \beta_i(t)g(S(t)) - \alpha_i(t)S(t). \qquad (17)
$$

In the special case of independent X_i, the total and class margins are always non-negative. Moreover, the natural allocation premium is always less than the stand-alone premium, meaning it satisfies the no-undercut condition of Denault (2001). Thus

$$
\mathsf{E}[X_i] \le \rho_X(X_i) \le \rho(X_i), \qquad (18)
$$

see Major and Mildenhall (2020) Proposition 2.

Returning to the general case, we show how eq. (17) gives useful insight into class margins. Since distortions are increasing and concave, $g(S(a)) \ge S(a)$ for all $a \ge 0$. Thus all asset layers contain a non-negative total margin. It is a different situation by class where

$$
M_i'(a) \ge 0 \iff \beta_i(a)g(S(a)) - \alpha_i(a)S(a) \ge 0 \iff \frac{\beta_i(a)}{\alpha_i(a)} \ge \frac{S(a)}{g(S(a))},
$$

so the class layer margin density is positive when β_i/α_i is greater than the average layer loss ratio $s/g(s)$. Since the loss ratio is ≤ 1 there must be a positive layer margin whenever $\beta_i(a)/\alpha_i(a) > 1$. But when $\beta_i(a)/\alpha_i(a) < 1$ it is possible the class has a negative margin. How can that occur depends on the relative thickness of the tails of X_i.

Defining an absolute notion of thin or thick tailed distribution is a very subtle matter. Embrechts, Klüppelberg, and Mikosch (1997) offer several different definitions. Fortunately, we are only concerned with relative tail thickness. We say that X_0 is **relatively thin tailed** or has **relatively thinner tails** than X_1 if $\kappa_0(x)$ is bounded as a function of x. Log-concave densities are all relatively thinner tailed than exponential and sub-exponential distributions. However, since $\sum_i \kappa_i(x) = x$ it is impossible for both X_0 and X_1 to be relatively thinner tailed. For a multivariate normal distribution, which has a log-concave density, the κ_i functions are all linear, showing it is not necessary that either class be relatively thin compared to the other. However, such balance is unusual.

Differentiating $\alpha_i(x)S(x) = \mathsf{E}[(X_i/X)1_{X>t}]$ with respect to x and re-arranging gives

$$\alpha_i'(x) = \left(\alpha_i(x) - \frac{\kappa_i(x)}{x}\right)\frac{f(x)}{S(x)}. \tag{19}$$

The results for β_i are analogous. The term $h(x) := f(x)/S(x)$ is called the hazard rate function. For thick tailed distributions it is an eventually decreasing function, but remains strictly positive. For thin tailed distributions it is an eventually increasing function. It is constant for the exponential distribution.

Equation (19) implies that $\alpha_i'(x) < 0$ if $\kappa_i(x)/x$ is decreasing (for example, if $\kappa_i(x)$ is bounded) because $\alpha_i(x)$ is the probability weighted integral of $\kappa_i(t)/t$ over $t > x$, and so $\alpha_i(x) < \kappa_i(x)/x$. Conversely if $\kappa_i(x)/x$ is increasing $\alpha'(i)$ will be positive.

In particular, if $\kappa_i(x)$ is bounded or if $\kappa_i(x)/x$ is decreasing then $\alpha_i(x)$ will decrease with x. Therefore $\beta_i(x) < \alpha_i(x)$ because the risk adjustment defining β_i weights tail events more. This shows it is possible that $\beta_i(x)/\alpha_i(x) < g(S(x))/S(x)$.

Usually $\alpha_i(x)$ eventually becomes constant or regularly varying like $1/x$ and so $\beta_i(x)/\alpha_i(x)$ will increase with x. At the same time, the layer loss ratio decreases with x because g is concave. Thus the thinner class will eventually get a positive margin density. Whether or not the thinner class has a positive total margin depends on the particulars of the classes and the level of assets a. As $a \to \infty$ the margin becomes positive because the natural allocation has the no-undercut property, eq. (18). A negative total margin it is more likely for less well capitalized insurers, which makes sense because they have a lower overall dollar cost of capital, and less expensive insurance pricing, because the loss ratio will be higher. We conclude that a relatively thin tailed class can have a negative margin, especially for weak capital standards giving low asset levels.

Since $\sum_i \kappa_i(x) = x$ it follows that $\sum_i \kappa_i'(x) = 1$. It is typical for the class with the thickest tail to behave like $\kappa_i(x) \approx t - \sum_{j\ne i} \mathsf{E}[X_j]$ for large t. In fact, this can be taken as a definition of thickest tail. Then $\kappa_i'(x) = 1$ and the remaining $\kappa_j(x) \approx \mathsf{E}[X_j]$ are almost constant. In that case $\kappa_j(x)/x > \alpha_j(x)$ and so $\alpha_j'(x) < 0$ and $\alpha_i'(x) > 0$; since $\sum_i \alpha_i(x) = 1$,

$\sum_i \alpha_i'(x) = 0$. It is only possible for two classes to have $\kappa_i(x) = \mathsf{O}(x)$ when their tails are balanced. Two compound Poisson distributions with the same severity provide an example.

Classes where $\alpha_i(x)$ is ultimately increasing with x always have a positive margin. Increasing α_i implies $\beta_i(x)/\alpha_i(x) > 1 > g(S(x))/S(x)$ because the risk adjustment puts more weight on X_i/X for larger X and the risk adjustment $g'(S(X))$ is comonotonic with X. Note that α_i must increase at least one i. The functions κ_i and α_i are not necessarily monotone; in fact κ_i can exhibit quite bizarre behavior.

This analysis shows that a relatively thin tailed class loses from pooling. By pooling it increases the assets available to pay losses, but in default states it will capture less than its unlimited expected loss proportion of the additional assets it contributes. Equal priority acts to transfer wealth away from the thiner class to the thicker one. Put another way, default is likely to be *caused* by a large loss in the thicker tailed line, which will then claim a disproportionate share of assets in default states. As a result the thinner class must be *paid* a margin, through negative M_i, to compensate for its losses in default states.

Next, we consider the behavior of α and β for losses below the expected value. Realistic, balanced insurance portfolios have skewed loss distributions with a very low probability of having an extremely good year and a small loss outcome. Thus for x small relative to expected loss it is usual that $f(x) \approx 0$ and $S(x) \approx 1$. Equation (19) then shows that $\alpha_i'(x) = 0$. Similarly $\beta_i'(x) = 0$.

Finally note that when $S(x)$ is very close to 1 the total margin $g(S(x)) - S(x)$ will be very close to zero. Thus the sum of class margins will be close to zero. Therefore, at small losses either every class has a zero margin or some are positive and some negative. As we have seen, a relatively thin tailed class will have a negative margin and the thickest tailed class—determined by ultimate behavior of $\kappa_i(x)$—will have a positive margin.

To conclude, we have shown the following behaviors.

1. In total, the margin and the margin density are always non-negative, for all layers and asset levels.
2. For full coverage, $a = \infty$, the natural allocation premium is always less than the monoline (stand-alone) premium and contains a non-negative margin.
3. As $a \to \infty$ all classes have a non-negative overall margin.
4. A relatively thinner tailed class will have a negative margin and margin density for low asset layers. It can have a negative overall margin and is more likely to do so with cheaper insurance pricing or a weaker capital standard.
5. A class with $\kappa_i(x) = \mathsf{O}(x)$ will always have a non-negative margin and margin density for all layers.

Figure 2 illustrates the theory we have developed. We refer to the charts as (r, c) for row $r = 1, 2, 3, 4$ and column $c = 1, 2, 3$, starting at the top left. The horizontal axis shows the asset level in all charts except $(3, 3)$ and $(4, 3)$, where is shows probability. Blue represents the thin tailed class X_0, orange thick tailed X_1, and green total X. Dashed lines represent losses and solid lines premium (or risk adjusted) when both are shown on the same plot. Here is the key.

- $(1,1)$ shows density for X_0, X_1 and $X = X_0 + X_1$; the two classes are independent. X_0 is a thin tailed gamma distribution with a mean of 150 and coefficient of variation of 0.1, which approximates US personal auto. X_1 is lognormal with a mean of 50 and $\sigma = 1.25$, which is quite extreme. The investor distortion uses a proportional hazard transform $g(s) = s^{0.5}$, resulting in moderately expensive insurance.
- $(1,2)$: log density; comparing tail thickness.
- $(1,3)$: the bivariate log-density. This plot illustrates where (X_0, X_1) *lives*. The diagonal lines show $X = k$ for different k. These show that large values of X correspond to large values of X_1, with X_0 about average.
- $(2,1)$: the form of κ_i is clear from looking at $(1,3)$. κ_1 has a local maximum around $x = 160$. κ_2 is monotonically increasing.
- $(2,2)$: The α_i functions. For small x the expected proportion of losses less than the expected 75/25 mix by Jensen's inequality. As x increases X_1 dominates. The two functions sum to 1.
- $(2,3)$: The solid lines are β_i and the dashed lines α_i from $(2,2)$. Since α_1 decreases $\beta_1(x) \le \alpha_1(x)$. This leads to class 0 having a negative margin in low asset layers. Class 1 is the opposite.
- $(3,1)$: illustrates premium and margin determination by asset layer for class 0 using eq. (9) and eq. (15). For low asset layers $\alpha_1(x)S(x) > \beta_1(x)g(S(x))$ (dashed above solid) corresponding to a negative margin density. Beyond about $x = 100$ the lines reverse and the margin density is positive.
- $(4,1)$: shows the same thing for class 1. Since α_1 is increasing, $\beta_1(x) > \alpha_1(x)$ for all x and so all layers get a positive margin. The solid line $\beta g S$ is above the dashed αS line.
- $(3,2)$: shows the layer margin densities. For low asset layers premium is fully funded by loss with zero overall margin. The thick class pays a positive margin and the thin class a negative one, reflecting the benefit the thick class receives from pooling in low layers. The overall margin is always non-negative. Beyond $x = 100$ the thin class margin is also positive.
- $(4,2)$: the cumulative margin in premium by asset level. Total margin is zero in low *dollar-swapping* layers and then increases. It is always non-negative. The lines in this plot are the integrals from 0 to a of those in $(3,2)$.
- $(3,3)$: shows stand-alone loss $(1 - S(x), x) = (p, q(p))$ (dashed) and premium $(1 - g(S(x)), x) = (p, q(1 - g^{-1}(1 - p))$ (solid, shifted left) for each class and total. The margin is the shaded area between the two. Each set of three lines (solid or dashed) does not add up vertically because of diversification. The same distortion g is applied to each line.
- $(4,3)$: shows the natural allocation of loss and premium to each class. The total (green) is the same as $(3,3)$. For each line, dashed shows $(p, \mathsf{E}[X_i \mid X = q(p)])$, i.e., the expected recovery conditioned on total losses $X = q(p)$, and solid shows $(p, \mathsf{E}[X_i \mid X = q(1 - g^{-1}(1 - p))])$, i.e., the natural premium allocation. Here the solid and dashed lines do add up vertically: they are allocations of the total. Looking vertically above p the shaded areas show how the total margin at that loss level is allocated between lines. Class 0 mostly consumes assets at low layers, and the blue area is thicker for small p, corresponding to smaller total losses. For p close to 1, large

total losses, margin is dominated by class 1 and in fact class 0 gets a slight credit (dashed above solid), reflecting the fact that it is partially providing capacity. The change in shape of the shaded margin area for class 0 is particularly evident: since class 0 is hurt by pooling, it pays a lower overall margin.

Plots $(3, 3)$ and $(4, 3)$ explain why the thick class pays relatively more margin: its margin shape does not change when it is pooled with class 0. In $(3, 3)$ the green shaded area is essentially an upwards shift of the orange and the orange areas in $(3, 3)$ and $(3, 4)$ are essentially the same. This means that adding class 0 has virtually no impact on the shape of class 1; it is like adding a constant. This can also be seen in $(4, 3)$ where the blue region is almost a straight line.

4.2 Implications for Market Structure

Throughout this section the eight assumptions from section 2.5 hold. In addition, we assume there are two classes of insured, with independent aggregate loss distributions X_0 and X_1. X_0 will usually be the less risky of the two. The regulatory capital standard is given by p-VaR for some p close to 1. VaR is the most commonly used regulatory capital standard. It underlies Solvency II.

A homogeneous loss model combined with positive homogeneous risk-based capital and investor pricing formula functionals implies the economics of any pool only depend on its mix of business and not on its size. Premium, loss, risk, and capital all scale with volume; loss ratio and return on equity are size-independent. This is an important simplification.

Since monoline pools within a class all have the same mix, their economics are independent of size. Therefore we can merge all the monoline pools by class and assume that the merged pool includes all the risks not in multiline pools. As a result, our homogeneous loss model is the perfect laboratory in which to study the pooling problem because it puts all the focus on the composition of pools.

In almost all situations there can only be one multiline pool. We shall show below that the natural premium rate for each risk in a multiline pool varies with the proportion of each class in the pool. Two pools with different proportions of each class have different rates by class. As a result, each class would have a definite preference for one of the two pools, and so they could not be in equilibrium. When the proportions are the same, the two pools can be combined by homogeneity. Therefore we can assume there is only one multiline pool.

Combining these two observations, we see the market structure is determined by a single decision variable: the mix by class in the multiline pool. Specifically, let t, $0 \le t \le 1$, denote the proportion of class 1 risk in the pool. By homogeneity, losses from a proportion t of class i are tX_i. Therefore losses from a pool with a proportion t of class 1 risks and $1-t$ of class 0 are given by

$$X(t) := (1-t)X_0 + tX_1. \tag{20}$$

The notation is chosen so that $X(0) = X_0$ and $X(1) = X_1$.

Different pools provide different covers. Although the premiums vary by class for pools

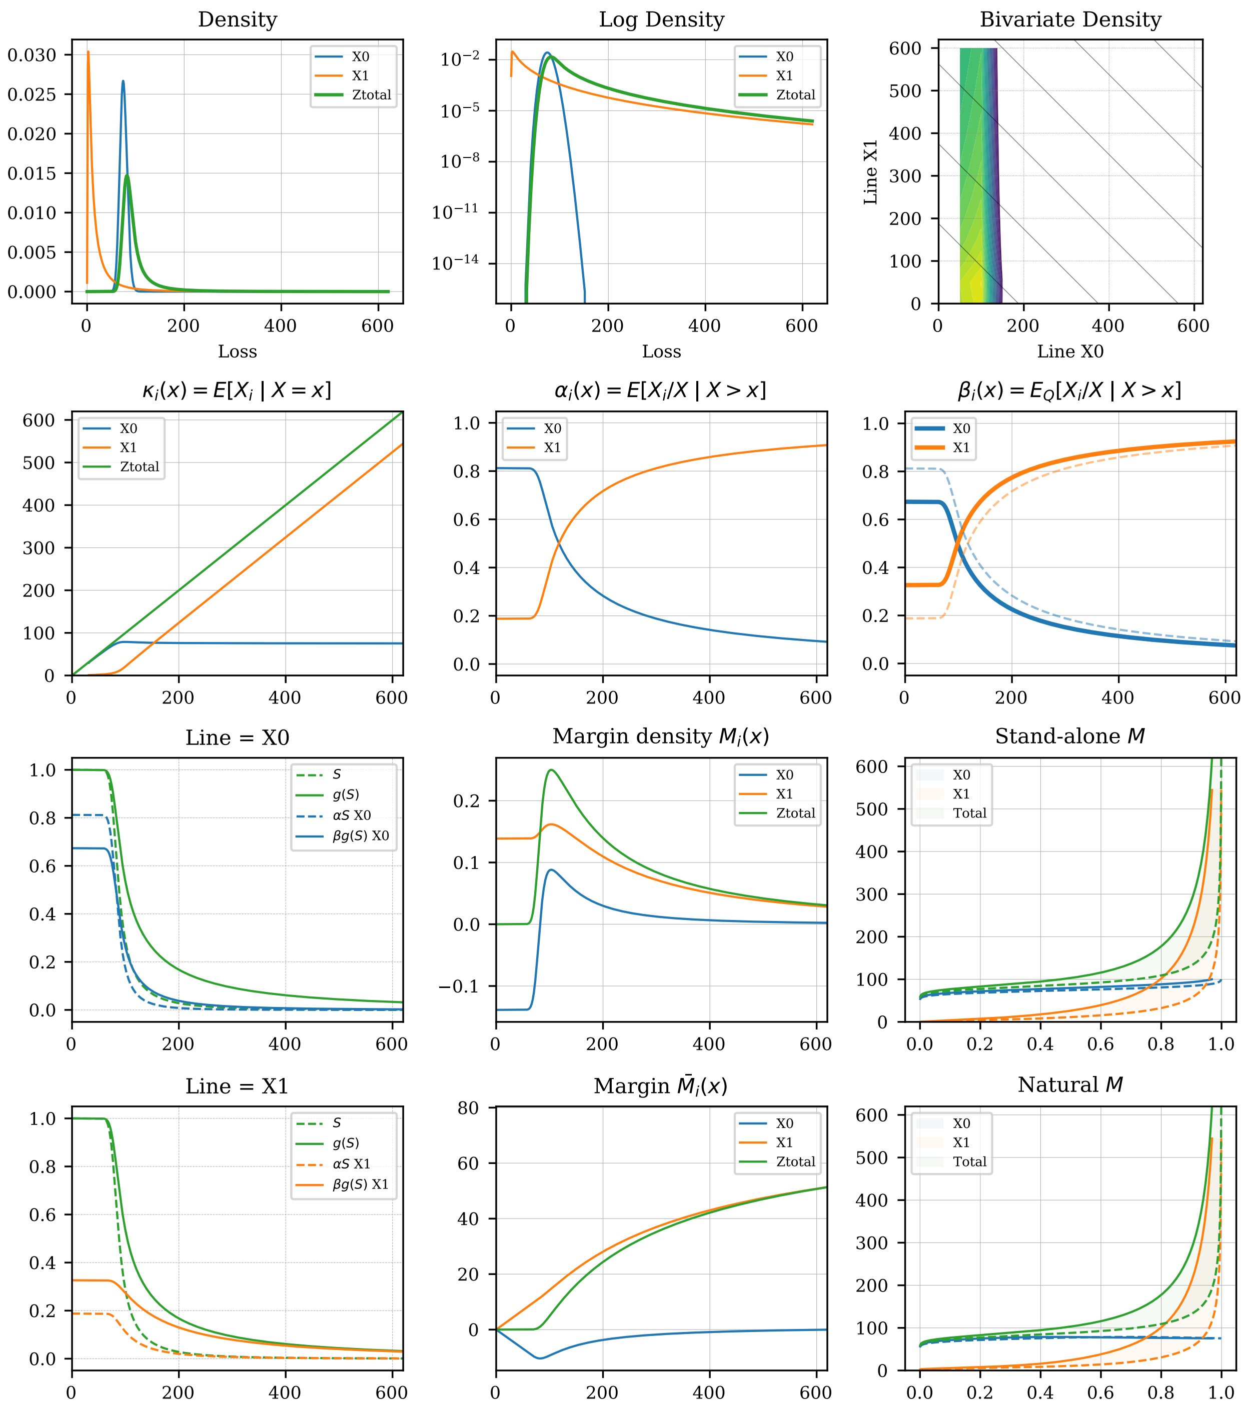

Figure 2: A thin tailed tailed class combined with a thick tailed line. See text for a key to the graphs.

with different mixes, all the premiums are fair values. Premium differences reflect different payments each pool will make in default states. Recall that insureds select between pools solely on the basis of price; they do not consider the payment differences.

If a multiline pool has $t \le 0.5$ then it can be scaled up to include all the class 0. If $t \ge 0.5$ it can scale to include all the class 1 risks. In both cases, homogeneity implies the economics are independent of pool scale. Thus there are just three structures that can occur:

1. $t = 0.5$: **complete pooling** in a single multiline pool;
2. $t = 0, 1$: **no pooling** with two disjoint monoline pools; or
3. $0 < t < 1$: $t \ne 0.5$: **partial pooling** with a multiline pool and a monoline pool.

In structure 3, all of class 0 is pooled for $0 < t < 0.5$ and all of class 1 for $0.5 < t < 1$. The remaining risks are monoline. Note that $t = 0$ and $t = 1$ define the same outcome. Outcomes are topologically a circle.

Our central question is to determine conditions under which each structure occurs. And when structure 3 occurs, what determines the proportions, whether an individual is pooled or monoline, and how is the diversification benefit distributed?

Let $P(t) := \rho(X(t) \wedge a(X(t)))$ be the premium for the pool $X(t)$. We denote the natural premium allocation and corresponding premium rates by

$$P_0(t) := \rho_{X(t)\wedge a(X(t))}((1-t)X_0), \quad P_1(t) := \rho_{X(t)\wedge a(X(t))}(tX); \tag{21}$$

$$R_0(t) := \frac{P_0(t)}{1-t}, \quad R_1(t) := \frac{P_1(t)}{t}. \tag{22}$$

There is no variable for the total pool rate because we have normalized the pool to contain one weighted policy, meaning its rate equals its premium. Since $P(t) = P_0(t) + P_1 t)$ we have $P(t) = (1-t)R_0(t) + tR_1(t)$. Note that $P(0) = P_0(0) = R_0(0)$: all three expressions equal the monoline premium for all of class 0. Similarly $P(1) = P_1(1) = R_1(1)$.

The most any insured will pay is their monoline premium, since they are free to form a single-policy monoline pool. Again, by homogeneity, the premium rate for a single policy pool is independent of its size.

Provided both monoline premiums are finite, $t \mapsto P(t)$ will be continuous; in fact, it is (generally) differentiable. Suppose the class distributions are continuous, so there are no issues with probability masses and the quantile functions are unique. Let S_t, f_t denote the survival and density of $X(t)$ and use the shorthand $a(t) = a(X(t))$. Then by eq. (6), combined with expressions for the derivatives of the survival and quantile function of homogeneous variables derived in Tasche (2001)[1], and the chain rule for differentiation, we get and expression for

[1]Given variables $X_1, \dots, X_n$, define a weighted homogeneous portfolio by $X(\mathbf{w}) = \sum_i w_i X_i$, $\mathbf{w} = (w_1, \dots, w_n)$. Let $S_{\mathbf{w}}(x) = \Pr(X(\mathbf{w}) > x)$ be the survival function, $q_{\mathbf{w}}(p)$ be the p quantile function, and $f_{\mathbf{w}}$ be the density of $X(\mathbf{w})$. Tasche obtains the following expressions.

The derivative of the survival function is

$$\frac{\partial S_{\mathbf{w}}}{\partial w_i}(x) = \mathsf{E}[X_i 1_{\{X(\mathbf{w})=x\}}] = \mathsf{E}[X_i \mid X(\mathbf{w}) = x] f_{\mathbf{w}}(x) = \frac{\kappa_i(x)}{w_i} f_{\mathbf{w}}(x). \tag{23}$$

the derivative of pool premium with respect to pool mix

$$
\begin{aligned}
\frac{dP}{dt} &= \frac{d}{dt}\int_0^{a(t)} g(S_t(x))\,dx \\
&= \int_0^{a(t)} g'(S_t(x))\frac{dS_t}{dt}(x)\,dx + g(S_t(a(t)))\frac{da}{dt} \\
+ &= \int_0^{a(t)} \left(\frac{\kappa_1(x)}{t} - \frac{\kappa_0(x)}{1-t}\right) g'(S_t(x)))f_t(x)\,dx + g(S_t(a(t)))\left(\frac{\kappa_1(a(t))}{t} - \frac{\kappa_0(a(t))}{1-t}\right) \\
&= \mathsf{E}_{\mathsf{Q}_t}[X_1 - X_0 \mid X_t \le a(t)](1-g(S_t(a(t))) + \mathsf{E}[X_1 - X_0 \mid X_t = a(t)]g(S_t(a(t)))
\end{aligned}
$$

where Q_t is the measure with density $g'(S_t(X_t))$. The second expectation in the last expression is the same its $\mathsf{E}_{\mathsf{Q}_t}$ version because DRMs are law invariant and therefore cannot change conditional probabilities within sets of the form $\{X_t = x\}$. This shows that the slope of P is controlled by κ_i.

By eq. (6), for $i = 0, 1$

$$
R_i(t) = \mathsf{E}_{\mathsf{Q}_t}[X_i \mid X_t \le a(t)](1-g(S_t(a(t))) + a(t)\mathsf{E}[X_i/X(t) \mid X_t > a(t)]g(S_t(a(t)))
$$

and therefore

$$
\begin{aligned}
(R_1(t) - R_0(t)) - \frac{dP}{dt} = a(t)\big(&\mathsf{E}_{\mathsf{Q}_t}[(X_1 - X_0)/X(t) \mid X_t > a(t)] - \\
&\mathsf{E}[(X_1 - X_0)/X(t) \mid X_t = a(t)]\big)g(S_t(a(t))). \qquad (26)
\end{aligned}
$$

Thus the slope of P is controlled by the difference between the conditional expectations on the right. Up to scaling, these are the difference between β and κ. The difference only appears in default states. As the capital standard gets stronger it disappears. It is also less material for thin tailed distributions. The result of conditioning on $\{X_t > a(t)\}$ and $\{X_t = a(t)\}$ will be more similar for a log-concave density than for thick tailed distributions. It also provides the missing link in our justification that there is only one multiline pool since $dP/dt \ne 0$ in most cases.

The derivative of $q_{\mathbf{w}}(p) = \mathsf{VaR}_p(X(\mathbf{w}))$ with respect to w_i is

$$
\frac{\partial q_{\mathbf{w}}(p)}{\partial w_i} = \mathsf{E}\left[X_i \mid X(\mathbf{w}) = q_{\mathbf{w}}(p)\right] = \frac{\kappa_i(x)}{w_i}. \qquad (24)
$$

For reference, the derivative of tail value at risk is

$$
\frac{\partial \mathsf{TVaR}_p(X(\mathbf{w}))}{\partial w_i} = \mathsf{E}[X_i \mid X(\mathbf{w}) > q_{\mathbf{w}}(p)] = \frac{\alpha_i(q_{\mathbf{w}}(p))}{w_i} \qquad (25)
$$

In our application, $X(t) = (1-t)X_0 + tX_1$, and so

$$
\frac{dS_t}{dt}(x) = -\frac{\partial S_t}{\partial w_0}(x) + \frac{\partial S_t}{\partial w_1}(x) = \left(\frac{\kappa_1(x)}{t} - \frac{\kappa_0(x)}{1-t}\right) f_t(x)
$$

and similarly for the quantile (VaR) derivative.

Consider adding a small amount of class 0 to a monoline portfolio of class 1. What are bounds on the multiline rates for each class? By definition $R_1(1) = P_1(1)$. What can we say about $R_0(t), R_1(t)$ as $t \to 1$ from below? Since the distributions are continuous $S_t(a(t)) = 1 - p$ for all t. Therefore, from eq. (6) and eq. (14)

$$R_0(t) = \mathsf{E}_{\mathsf{Q}_t}[X_0 | X(t) \le a](1 - g(1-p)) + (a/(1-t))\beta_0(a)g(1-p) \tag{27}$$

$$R_1(t) = \mathsf{E}_{\mathsf{Q}_t}[X_1 | X(t) \le a](1 - g(1-p)) \ + (a/t)\beta_1(a)g(1-p). \tag{28}$$

Suppose that class 1 is relatively thin tailed.

First consider t close to 1, meaning class 0 has a small volume. When class 0 is added in small quantities it acts like adding a constant to the loss. It will increase assets by $(1-t)\mathsf{E}[X_0]$ and will have a de minimus impact on total losses. In non-default states everything is fair. But in default states, which will be driven by a large class 1 loss, class 1 will take the lion's share of the additional assets under equal priority. There is no offsetting compensation to class 0 because it is relatively thin tailed: it never has an out-sized loss. Thus class 0 effectively subsidizes class 1 by providing more asset for class 1 in default. To compensate for this, class 1 must pay a higher premium: it is now better off than it was in a monoline, $t = 1$ pool. Thus $R_1(t) > R_1(0)$ for t close to 1. The amount by which the pooled rate exceeds the monoline rate will increase as the capital standard becomes weaker because default states become more important.

On the other hand for class 0 we can see that the economic value of recoveries will be close to $\mathsf{E}[X_0]$. It will certainly be less than $\mathsf{E}[X_0 g'(S(X))]$, which does not allow for default. Thus

$$\begin{aligned} R_0(t) &\le \mathsf{E}[X_0 g'(S(X))] \\ &= \mathsf{E}[X_0]\mathsf{E}[g'(S(X))] + \mathsf{Cov}(X_0, g'(S(X))) \\ &= \mathsf{E}[X_0] + \mathsf{Cov}(X_0, g'(S(X(t)))), \end{aligned}$$

and $\mathsf{Cov}(X_0, g'(S(X(t)))$ will be close to zero for t close to 1 and can in fact be negative in certain situations. Conclusion 4 in section 4.1 implies $R_0(t) < \mathsf{E}[X_0]$ (actually $< \mathsf{E}[X_0 \wedge a(0)]$). If class 0 is not relatively thinner tailed then $R_0(t) \ge \mathsf{E}[X_0]$ by conclusion 5.

Second, consider the opposite case: adding a small amount of a more volatile portfolio to a less volatile one. This corresponds to t close to 0. This situation is different to the previous one because if X_1 has thick enough tails then it can still dominate large outcomes even when it has a very small expected loss. The conclusions in section 4.1 show the relationship between $R_0(t)$ and $P(0)$ is indeterminate; it depends on the relative tail thicknesses. When class 1 is not extremely thick tailed $R_0(t)$ will exceed $P(0)$ for small g, but the effect is much less pronounced than the situation at $t = 1$.

We conclude that if class 0 is relatively thinner tailed than class 1 then $R_1(t) > P(1)$ for t in an interval $[t_1^*, 1)$. When class 0 is not relatively thinner then $R_1(t) > \mathsf{E}[X_1]$ and there can exist $t_1^* < t_1^{**}$ so that $R_1(t) \ge P(1)$ for $t \in [t_1^*, t_1^{**}]$ but $\mathsf{E}[X_1] \le R_1(t) \le P(1)$ for $t \in [t_1^{**}, 1)$.

Figure 3 illustrates these variables for a typical example. X_0 and X_1 are independent gamma variables with $\mathsf{E}[X_i] = 100$ and coefficient of variation 0.25 and 0.30 respectively. Both lines are thin tailed and have log concave densities; neither is relatively thinner than the other. To show various features more clearly, the model uses a weak 0.90 VaR capital standard and an expensive proportional hazard ρ with distortion $g(s) = s^{0.3}$.

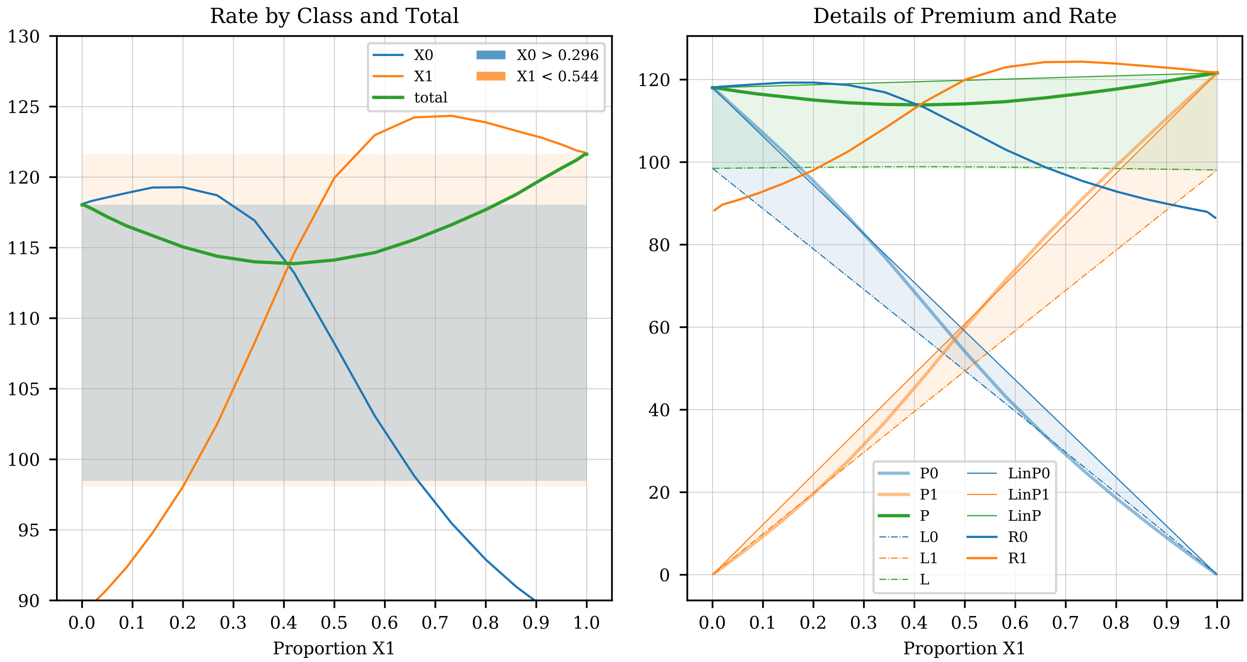

Figure 3: Typical behavior of pricing variables. Horizontal axis t determines the pool mix as $(1-t)X_0 + tX_1$.

The plot on the left shows:

- The premium rate $R_i(t)$ for each class (blue and orange) and total premium $P(t)$ (green). The three lines intersect at a point: this will always occur because the total premium is a $(1-t, t)$ weighted average of the by-line rates. They need not intersect at the minimum.
- The horizontal shaded bands show the range from monoline expected loss up to the monoline premium for each class. Because the capital standard is weak monoline expected losses are $< \mathsf{E}[X]$ by limited liability. Since X_1 is more risky than X_0 it has a higher monoline premium.
- As discussed above, $R_1(1) = P(1)$ and it bows up for t close to 1. Here $t_1^* = 0.544$. This reflects the fact it benefits from pooling with X_0. Similar comments apply to R_0 near $t = 0$; its rate is above the monoline rate until $t_0^* = 0.296$. The critical t values are shown in the legend.
- The total premium line only gives the premium for the $(1-t, t)$ weighted pool; it does not give the total market premium because the blended portfolio can only include all risks in the special case $t = 0.5$. When $t \neq 0.5$ the multiline pool with include all of one class, but the other will be split and partially written in a monoline pool.

By market assumption 1, that insureds are pure price buyers, a market with prices greater than monoline rates will not be in equilibrium: risks will defect into a cheaper monoline pool.

Thus the feasible range for a solution is $t_0^* < t < t_1^*$. For $t < t_0^*$ the pool rate for class 0 is greater than the monoline rate, so it will not pool. And for $t > t_1^*$ the same holds for class 1.

We can now determine the market equilibrium.

- If $t_0^* < 0.5 < t_1^*$ the we are in structure 1 and have complete pooling. If $t \neq 0.5$ then risks left in the monoline pool, paying a higher rate, will offer to pool at a slightly more advantageous proportion, and lower rate, for the other class and the pool unravels. The pool is only in equilibrium when there are no risks left in monoline pools, i.e., when $t = 0.5$. Examples show this is a very common outcome, especially when the two classes are of comparable volatility.
- If $t_0^* > t_1^*$ then we are in structure 2 with no pooling and two monoline pools. The feasible region only contains the point $t = 0 = 1$, remember 0 and 1 define the same outcome. This structure occurs for example when X_0, X_1 are very thick tailed and VaR is not subadditive.
- If $t_0^* < t_1^* < 0.5$ then the equilibrium is at $t = t_1^*$ and we have partial pooling, structure 3. The multiline pool contains all of class 0 and a proportion of class 1. Class 0 gets the benefits of diversification and pays less than its monoline rate. Class 1 pays its monoline rate. Class 0 has the negotiating power: it only needs to attract a portion of the class 1 risks and it knows there are always class 1 risks paying the monoline rate, so that acts as a bogey. It can offer to pool at a proportion $t_1^* - \epsilon$ for small $\epsilon > 0$. This produces a price $R_1(t_1^* - \epsilon) < R_1(t_1^*) = R_2(1)$ below the monoline price, which will be enough to attract as many class 1 risks as needed. If $\epsilon > 0$ then the class 1 risks who remain in a monoline pool, paying a higher rate, have an incentive to offer to pool at a share $t_1^* - \epsilon/2$ to class 0. The original pool will unravel as this will be a lower price for class 0. Thus the equilibrium is at t_1^*. Class 1 has no negotiating power because there will always be risks in its monoline pool, whereas all class 0 is entirely in the multiline pool. Similarly if $0.5 < t_0^* < t_1^*$ then the multiline pool contains all class 1 and some of class 0. Now class 0 pays the monoline rate and class 1 benefits from the pooling with a lower rate. The equilibrium is at $t = t_0^*$.

All the equilibriums are Pareto optimal because of the shape of $R_0(t)$ and $R_1(t)$ at t: one is increasing and one decreasing so it is not possible to make both classes better off. Changing t will increase the rate for at least one class.

Returning to fig. 3 we see that its equilibrium is complete pooling, structure 1, since $t_0^* = 0.206 < 0.5 < 0.544 = t_1^*$.

It is highly unlikely the solution is at the minimum of $P(t)$. That solution only includes all risks when the minimum occurs at $t = 0.5$.

If we introduce reinsurance then it can be used in structure 3 to provide a blended rate to the split class. Suppose class 1 is split. The reinsurer enables a solution with proportion t where $t_0^* < t < t_1^*$ by assuming class 1 from the monoline and the pool in return for a blended rate. However, without an enforcement mechanism this is not a stable solution. Members of class 0 have an incentive to offer reinsurance pool participants the true, lower, multiline rate unless $t = t_1^*$. In that case the pool would unravel as before.

As the capital standard strengthens the rates bow up less and the feasible region gets broader. With unlimited liability, i.e., $a = \infty$, the market pricing functional reduces to the investor functional, which is coherent. In that case pooling is always optimal provided both premiums are finite. This follows from Proposition 2: for a coherent risk measure the natural allocation always contains a non-negative margin and is less than the monoline rate (no-undercut). Thus the natural allocation premium $P_1(t)$ satisfies $\mathsf{E}[tX_1] \le P_1(t) \le \rho(tX_1)$. Since ρ is positive homogeneous, dividing by t gives $\mathsf{E}[X_1] \le R_1(t) \le \rho(X_1)$. Note $\rho(X_1)$ is the monoline premium for unlimited cover.

The right hand plot of fig. 3 illustrates these points. The two shaded triangles show the segments bounded by $\mathsf{E}[X_i]$ and $\rho(X_i)$. The thicker premium line is within the segment precisely when the rate is between monoline loss cost and premium. The green band shows the corresponding bounds for total premium.

The following behaviors are evident.

1. Full pooling is more likely with a stronger capital standard. In the limit of full coverage, we always get full poling provided the full coverage premiums are both finite. If one or other premium is infinite full pooling will not be an equilibrium solution. Coherent risk measures on unbounded variables must sometimes assume infinite values by Delbaen (2002), Theorem 5.1.
2. For thin tailed lines with log concave densities full pooling is more likely with more expensive pricing. (Expensive pricing means a greater markup over expected loss.) The opposite is true for thick tailed lines. Thin tailed lines are concentrated near the mean and large losses occur when each component has a somewhat above average loss. In this situation pooling is not greatly beneficial. thick tailed lines are concentrated away from the mean. A large loss occurs from a large loss on one component and a small loss on the other. This situation benefits from pooling and results in superior pooled coverage, which therefore costs more. Thus the market pricing operator fails to be subadditive.
3. Classes with unbalanced tails are more likely to result in solution 3 than balanced classes. This follow for the same reason as the second case of 2.

5 Examples

5.1 Structure 2: No Pooling

In fig. 4 X_0 is a lognormal with unlimited mean 100 and $\sigma = 1.5$ and X_1 is a very heavy tailed Pareto with unlimited mean 10 and $\alpha = 1.2$. The right hand plot shows the (green) premium line bows up from the line connecting the two monoline rates, meaning there can be no pooled solution. The weighted average of two rates lower than the monoline rates has to lie below the line connecting them.

In this example $P(t)$ has a maximum at $t < 1$ and is actually decreasing for t close to 1.

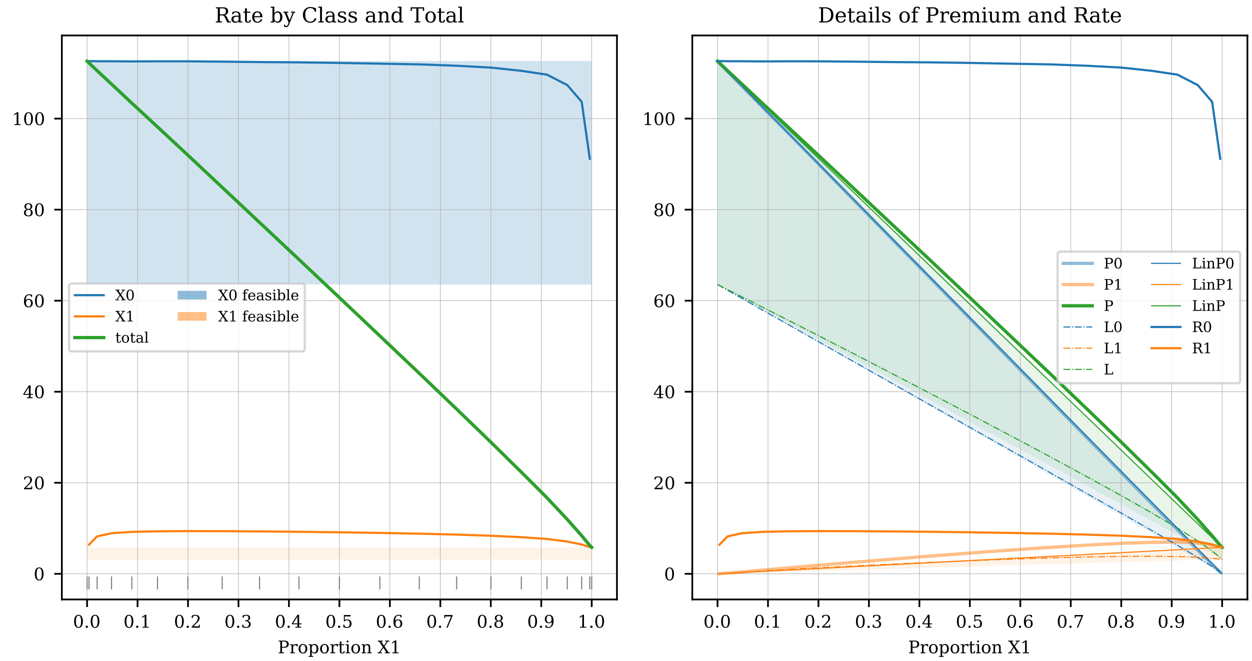

Figure 4: Market equilibrium structure 2: two monoline pools.

Pooling a small amount of (the much larger) X_0 greatly increases the assets available and X_1 captures an outsize proportion of them in default because of its thick tail. As a result its absolute premium increases as t decreases from 1, not just its premium rate.

Since X_0 is actually quite thick tailed $R_0(t) > \mathsf{E}[X_0]$ for all t: the blue rate line stays in the shaded blue area on the left.

5.2 Structure 3: Partial Pooling

In fig. 6 illustrates several interesting features. X_0 is a lognormal with $\sigma = 1.5$ and unlimited mean 150; X_1 is Pareto with unlimited mean 50 and $\alpha = 1.9$, so both lines have very thick tails. The capital standard is VaR 0.9 and pricing uses a proportional hazard $g(s) = s^{1/4}$.

The figure shows the following features.

1. $P_1(t) > P(1)$ for t close to 1.
2. P has an inflexion point and is sub-additive for $t \in [0, 0.7]$ and super-additive for $t \in [0.7, 1.0]$.
3. The feasible region is very small, approximately $[0.628, 0, 647]$, which does not include $t = 0.5$ and therefore we have partial pooling, structure 2.

The drop in R_1 for t close to zero is a numerical artifact.

5.3 A Realistic Example

To draw out different behaviors, the two previous examples use more extreme parameters that would normally apply to an insurance book. Our final example shows more typical

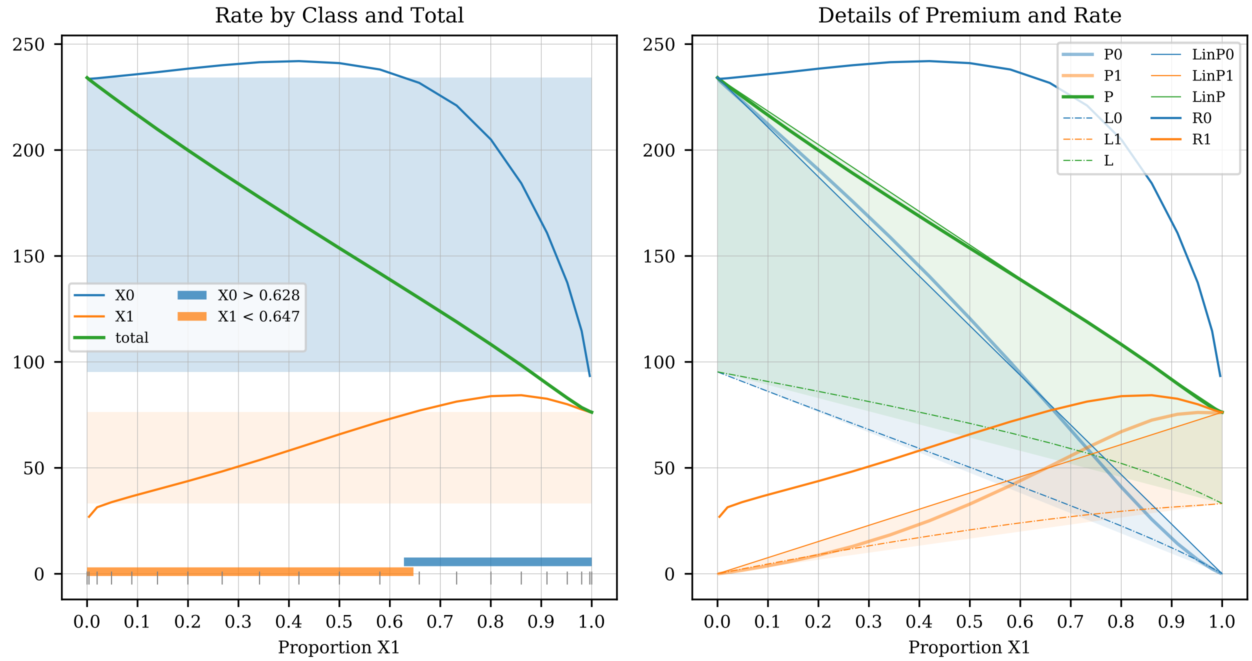

Figure 5: Market equilibrium structure 3: partial pooling.

parameter values to illustrate a generic view. It results in a full pooling structure. Class 0 is gamma, mean 150 and coefficient of variation 0.15 typical for a low limit liability book. Class 1 has a lognormal distribution with a mean of 100 and $\sigma = 0.3$. It represents a higher limit book. The distortion is $g(s) = s^{0.8}$, which is in-line with distortions calibrated to market pricing, and the capital standard is Sovlvency II, $p = 0.995$.

It may appear that fig. 6 shows only small differences between classes. Low volatility insurance is a very competitive business, written with very thin margins. Figure 7 shows the implied loss ratios by class. These are comparable with combined ratios, since our model excludes expenses. The loss ratio differences are material. The equilibrium, $t = 0.5$ shows class 0 written at 98.4 percent vs. its monoline rate of 97 percent. Class 1 is written at 94.5 percent vs. 93.5 percent. Thus class 0 achieves a 1.4 percentage point decrease in loss ratio from pooling vs. 1 point for class 1.

5.4 Note on the Computations

The computations underlying each figure were performed using discrete approximations with 2^{16} equally sized buckets and a sample of 21 values of t in $[0, 1]$ inclusive. Convolutions are performed using Fast Fourier Transforms (FFT), Grubel and Hermesmeier (1999), Mildenhall (2005). The calculations are essentially exact other than a minor discretization error. The conditional expectations needed for κ_i are also be performed using a FFT

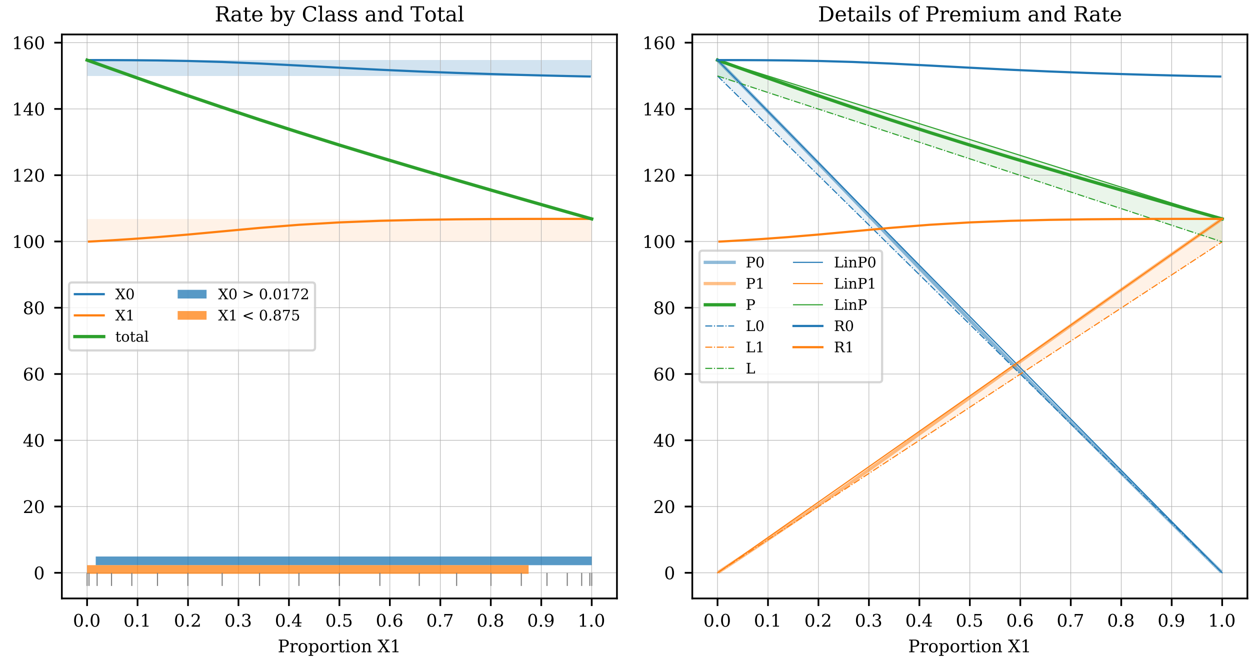

Figure 6: Market equilibrium structure 1: full pooling with parameters reflecting a liability and property insurance portfolio.

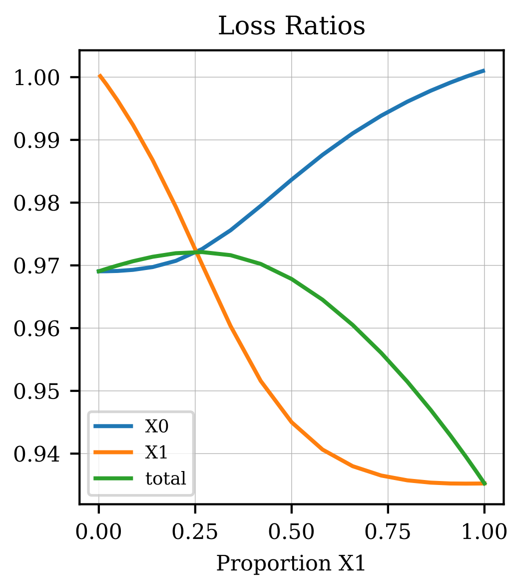

Figure 7: Loss ratios by class across different portfolio corresponding to fig. 6.

6 Conclusions

We have presented a novel but realistic model of a two class insurance market. The market includes a compulsory insurance requirement, capital regulation, and costly capital but otherwise is efficient. Depending on the aggregate loss characteristics of the two risk classes the Pareto optimal market equilibrium can be two monoline pools, a monoline and a multiline pool, or one multiline pool. In general when the classes are comparably risky there is one multiline pool. When a more risky class is combined with a less risky one, the less risky class often gets the benefit of pooling and pays a rate below its monoline premium while the more risky class pays its monoline premium. Stricter capital standards make more complete pooling more likely because it increases the importance of economizing on capital. There is no pooling when the risk have extremely thick tails and the capital standard is not subadditive.

Reinsurance can be used to pool premium rates, but without an enforcement mechanism it does not provide a stable solution.

The results are consistent with observed market structure in US property-casualty insurance, where more volatile lines are often written by monoline companies. Florida homeowners and medical malpractice liability are two examples. It is also consistent with the existence of highly leveraged, low risk pools, such as monoline auto writers.

References

Acerbi, Carlo, 2002, Spectral measures of risk: A coherent representation of subjective risk aversion, *Journal of Banking & Finance* 26, 1505–1518.

Anscombe, F. J., and R. J. Aumann, 1963, A Definition of Subjective Probability, *The Annals of Mathematical Statistics* 34, 199–205.

Aon Benfield, 2015, *Insurance Risk Study*. 10th ed.

Arrow, Kenneth J., 1996, The theory of risk-bearing: Small and great risks, *Journal of Risk and Uncertainty* 12, 103–111.

Boonen, Tim J., Andreas Tsanakas, and Mario V. Wuthrich, 2016, Capital allocation for insurance portfolios with non-linear risk aggregation,, 1–31.

Campi, Luciano, Elyès Jouini, and Vincent Porte, 2013, Efficient portfolios in financial markets with proportional transaction costs, *Mathematics and Financial Economics* 7, 281–304.

Castagnoli, Erio, Fabio Maccheroni, and Massimo Marinacci, 2002, Insurance premia consistent with the market, *Insurance: Mathematics and Economics* 31, 267–284.

Castagnoli, Erio, Fabio Maccheroni, and Massimo Marinacci, 2004, Choquet insurance pricing: A caveat, *Mathematical Finance* 14, 481–485.

Chateauneuf, A., R. Kast, and A. Lapied, 1996, Choquet pricing for financial markets with frictions, *Mathematical Finance* 6, 323–330.

Cummins, J David, 1988, Risk-Based Premiums for Insurance Guarantee Funds, *Journal of Finance* 43, 823–839.

Cummins, J. David, 2000, Allocation of capital in the insurance industry, *Risk Management and Insurance Review*, 7–28.

Cummins, J. David, Georges Dionne, and Robert Gagne, 2008, The Costs and Benefits of Reinsurance,.

Cummins, J. David, Mary A. Weiss, Xiaoying Xie, and Hongmin Zi, 2010, Economies of scope in financial services: A DEA efficiency analysis of the US insurance industry, *Journal of Banking and Finance* 34, 1525–1539.

Delbaen, Freddy, 2000, Coherent risk measures (Pisa Notes), *Blätter der DGVFM* 24, 733–739.

Delbaen, Freddy, 2002, Coherent risk measures on general probability spaces, *Advances in Finance and Stochastics: Essays in Honour of Dieter Sondermann*, 1–37.

Denault, Michel, 2001, Coherent allocation of risk capital, *Journal of risk*, 1–39.

Denneberg, Dieter, 1994, *Non-additive measure and integral* (Kluwer Academic, Dordrecht).

De Waegenaere, Anja De, 2000, Arbitrage and Viability in Insurance Markets, *The GENEVA Papers on Risk and Insurance Theory* 25, 81——99.

De Waegenaere, Anja, Robert Kast, and Andre Lapied, 2003, Choquet pricing and equilibrium, *Insurance Mathematics and Economics2* 32, 359–370.

Dietz, Simon, and Oliver Walker, 2017, Ambiguity and Insurance: Capital Requirements and Premiums, *Journal of Risk and Insurance* 86, 213–235.

Doherty, Neil A, and S. M. Tinic, 1981, Reinsurance under conditions of capital market equilibrium: a note, *Journal of Finance*.

Eeckhoudt, Louis, Christian Gollier, and Harris Schlesinger, 2011, *Economic and financial decisions under risk* (Princeton University Press).

Ellsberg, Daniel, 1961, Risk, Ambiguity, and the Savage Axioms, *The Quarterly Journal of Economics* 75, 643–669.

Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch, 1997, *Modelling Extremal Events* (Springer Verlag, Berlin Heidelberg).

Föllmer, Hans, and Alexander Schied, 2011, *Stochastic finance: an introduction in discrete time.* third edit. (Walter de Gruyter).

Grubel, Rudolf, and Renate Hermesmeier, 1999, Computation of Compound Distributions I: Aliasing Errors and Exponential Tilting, *Astin Bulletin* 29, 197–214.

Grundl, Helmut, and Hato Schmeiser, 2007, Capital allocation for insurance companies—What Good Is It?, *Journal of Risk and Insurance* 74.

Huber, Peter J., 1981, *Robust statistics* (John Wiley & Sons, Inc.).

Ibragimov, Rustam, Dwight Jaffee, and Johan Walden, 2018, Equilibrium with monoline and multiline structures, *Review of Finance* 22, 595–632.

Ibragimov, Rustam, Dwight Jaffee, and Johan Walden, 2010, Pricing and Capital Allocation for Multiline Insurance Firms, *Journal of Risk and Insurance* 77, 551–578.

Ibragimov, Rustam, and Johan Walden, 2007, The limits of diversification when losses may be large, *Journal of Banking and Finance* 31, 2551–2569.

Jiang, Wenjun, Marcos Escobar-Anel, and Jiandong Ren, 2020, Optimal insurance contracts under distortion risk measures with ambiguity aversion, *ASTIN Bulletin* 50, 1–28.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, A smooth model of decision making under ambiguity, *Econometrica* 73, 1849–1892.

Kusuoka, Shigeo, 2001, On law invariant coherent risk measures, *Advances in Mathematical Economics* 3, 83–95.

Major, John A, and Stephen J Mildenhall, 2020, Charging for Diversifiable Risk and Proud to Do It: Multiline Insurance Pricing With a Distortion Risk Measure,.

Marshall, Daniel, 2018, An Overview of the California Earthquake Authority, *Risk Management and Insurance Review* 21, 73–116.

Mildenhall, Stephen J, 2005, Correlation and Aggregate Loss Distributions With An Emphasis On The Iman-Conover Method, *Casualty Actuarial Society Forum, WinterSociety.*

Mildenhall, Stephen J, 2017, The Coming Revolution in Actuarial Modeling—Election Day Lessons for the Predictive Data Analyst, *Actuarial Review*, 38——43.

Myers, Stewart C, and James A Jr. Read, 2001, Capital allocation for insurance companies, *Journal of Risk and Insurance* 68, 545–580.

Phillips, Richard D., J. David Cummins, and Franklin Allen, 1998, Financial Pricing of Insurance in the Multiple-Line Insurance Company, *Journal of Risk and Insurance* 65, 597–636.

Robert, Christian Y., and Pierre-E. Therond, 2014, Distortion Risk Measures, Ambiguity Aversion and Optimal Effort, *ASTIN Bulletin* 44, 277–302.

Schmeidler, David, 1986, Integral representation without additivity, *Proceedings of the American Mathematical Society* 97, 255–255.

Schmeidler, David, 1989, Subjective Probability and Expected Utility without Additivity, *Econometrica* 57, 571–587.

Shapiro, Alexander, 2012, On Kusuoka Representation of Law Invariant Risk Measures, *Mathematics of Operations Research* 38, 142–152.

Sherris, Michael, 2006, Solvency, capital allocation, and fair rate of return in insurance, *Journal of Risk and Insurance* 73, 71–96.

Svindland, Gregor, 2009, Continuity properties of law-invariant (quasi-)convex risk functions on L∞, *Mathematics and Financial Economics* 3, 39–43.

Tasche, Dirk, 2001, Conditional expectation as quantile derivative, *Arxiv preprint math/0104190*, 1–12.

Tsanakas, Andreas, and Christopher Barnett, 2003, Risk capital allocation and cooperative pricing of insurance liabilities, *Insurance: Mathematics and Economics* 33, 239–254.

Venter, Gary G., John A. Major, and Rodney E. Kreps, 2006, Marginal Decomposition of Risk Measures, *ASTIN Bulletin* 36, 375–413.

Wang, Shaun, 1996, Premium Calculation by Transforming the Layer Premium Density, *ASTIN Bulletin* 26, 71–92.

Wang, Shaun, Virginia Young, and Harry Panjer, 1997, Axiomatic characterization of insurance prices, *Insurance: Mathematics and Economics* 21, 173–183.

Yaari, Menahem E., 1987, The Dual Theory of Choice under Risk, *The Econometric Society* 55, 95–115.

Zhang, Jiankang, 2002, Subjective Ambiguity, Expected Utility and Choquet Expected Utility, *Economic Theory* 20, 159–181.