BETA
Bureau
d'économie
théorique
et appliquée

Documents de travail

« Risk pooling and ruin probability, or why high risks are not bad risks »

Auteurs

Debora ZAPAROVA, Sandrine SPAETER

Document de Travail n° 2019 – 33

Septembre 2019

Bureau d'Économie Théorique et Appliquée BETA

www.beta-umr7522.fr

@beta_economics

Contact :
jaoulgrammare@beta-cnrs.unistra.fr

Risk pooling and ruin probability, or why high risks are not bad risks

Debora Zaparova, Sandrine Spaeter
BETA, University of Strasbourg

September, 2019

Abstract

The aim of this paper is to show that high risks being bad risks is a misinterpretation. We discuss the role of the risk-bearing capital in the insurance process. In particular, we explicit the link between insurer's risk, capital, size and composition of an insurance pool. Those parameters have a tangible impact on the insurer's ruin probability when his size is limited. An additional policyholder may increase the ruin probability, while a specific combination of risk types may produce a significant decrease. Those implications should be considered given the legal requirements relative to the insurer's insolvency. A strategy that consists in attracting only low-risk agents is not necessarily expedient for an insurer.

Keywords: high risks, insurance, risk loading

JEL classification: D81, G22, G28

1 Introduction

Technological advance produces new means of data generation, new ways of gaining information and new tools for data analysis, all of which may be used to evaluate individual risk profiles and decrease the uncertainty on risk parameters. It also produces various misgivings about the use of information. One of the public concerns is insurance cream-skimming. The ability to discover a policyholder's individual risk level might provide an incentive for insurers to attract only "good" agents with a low individual probability of becoming claimants.

In this paper, we want to point out that high risks being bad risks is a misinterpretation. High-risk policyholders contribute to the insurer's capacity to cover losses. Consequently, a strategy that consists in attracting only low-risk agents is not necessarily expedient for an insurer. In order to develop our argument, we examine the role of the risk-bearing capital in the insurance process. In particular, we explicit the link between insurer's risk, capital, size and composition of an insurance pool.

Risk pooling allows to reduce possible variations of one's wealth by subdividing the global variance of the pool among all the participants (Borch, 1962). Yet, the risk pooling is not only linked to the individual risk, but also to the insurer's risk, which is the risk of ruin.

The risk of ruin is the probability that the insurer will not be able to cover all the losses experienced by the policyholders. Such a possibility is disregarded in the economic literature. This is due to the assumption that an insurer can create an infinitely large pool. In this case, the average loss per individual converges to its expected value. Yet, the size of the pool per se does not imply that the risk of ruin converges to zero. The decrease in the ruin probability is not due to the presence of numerous policyholders, but to the capital that they provide.

The risk of ruin decreases with the number of policyholders only if each of them contributes an additional amount on top of the pure premium equal to the individual expected loss. An increase in the pool size leads to an increase in the global expected loss. But it also produces a more than proportional increase of the financial reserves, due to those additional contributions which are collected through the premium loading[1]. The loaded premiums form a risk-bearing capital which acts as a buffer fund used to cover the claims above the expected average. Thus, the assumption of an infinitely large insurance pool is useful in order to state the following: if each policyholder contributes an extra amount to the insurer's buffer fund, it makes an aggregate funds insufficiency extremely unlikely.

The role of the risk loading and the equity capital in the insurance process is seemingly neglected in the economic literature. The discussion on the role of the buffer fund in the reduction of insurer's risk can be found in Houston (1964), Cummins (1974), Venezian (1984) and Cummins (1991). Cummins (1974) and Smith and Kane (1994) also mention various misinterpretations of the risk pooling implications

[1] Such a loading is called a safety or a risk loading.

for the insurer. Furthermore, Smith and Kane (1994) discuss the fact that the ruin probability is discontinuous and point out the importance of this observation for the small pools. More recent literature on the role of the risk loading and its relation to the ruin probability includes Powers, Venezian, and Jucá (2003), Powers (2006), Denuit, Eeckhoudt, and Menegatti (2011), Gatzert and Schmeiser (2012) and Klein and Schmeiser (2018).

In this paper, we suggest that it is not necessarily desirable from the insurer's point of view to form an homogeneous low-risk pool. As Feller (1968) shows in case of Bernoulli trials, under some conditions the heterogeneity contributes to the reduction of the global risk in the pool[2]. We explicit the link between ruin probability, capital, size and composition of an insurance pool. Those parameters have a tangible impact on the insurer's risk when the number of potential policyholders is limited. For instance, an additional policyholder may increase the ruin probability, thus questioning the idea that having more policyholders is always better for the risk pooling process. We also examine the role of the high-risk agents and show that high-risk policyholders contribute to the accumulation of the risk-bearing capital through the risk loading. Moreover, the heterogeneity implies that it is possible to define a specific risk type combination which may significantly decrease the risk of ruin.

We argue that insurers should take these implications into account in order to manage their risk of insolvency and to keep risk pooling attractive for the policyholders. In particular, European insurers have to reduce and control their ruin probability according to the legal terms of policyholders protection, such as Solvency II. Solvency II is the European Union Directive that came into effect on January 2016[3]. Its mission is to regulate the European insurance market, in order to improve the policyholder protection and enhance the market harmonization. An insurer has to correctly evaluate his risk and to form a proper buffer fund in order to maintain his ruin probability under a required level of 0.5% a year. He might do it by creating a pool large enough to provide the amount of capital that is necessary. Or he might raise the premiums in order to compensate for the pool size. Since an homogeneous low-risk pool can be difficult to create, attracting only low-risk policyholders is not an expedient strategy.

Our paper is organized as follows. In Section 2, we discuss the risk-bearing capital

[2]In case of Bernoulli trials, when the expected loss is fixed for both pools, an homogeneous pool actually maximizes the global risk.

[3]Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II) (recast) (Text with EEA relevance).

and how it is affected by the size and composition of an insurance pool. We show that multiple combinations of policyholders allow to accumulate a given amount of risk-bearing capital, and thus to achieve a given level of risk-bearing capacity, which is not necessarily improved by an entry of an additional policyholder. In Section 3, we extend those observations to the insurer's risk of ruin, by showing that the ruin probability is reduced only due to the improvement of the risk-bearing capacity. We show that each discontinuity point of the ruin probability function is linked to a specific combination of policyholders. In Section 4, we discuss the implications of those results. We conclude in Section 5.

2 Risk-bearing capital, size and composition

Consider a population of N agents, each facing a risk of losing a monetary amount L. Individual risks are independent. Each agent can either be a high-risk type (h) or a low-risk type (l), with individual loss probability p_k, $k = h, l$. High-risk agents are characterized by a higher loss probability $p_h > p_l$. Thus, an individual risk of loss is a random variable X_k with an expected value $\mathbb{E}(X_k) = p_k L$.

Each agent can cover his risk of loss by signing a contract with an insurance company. The size of the insurance pool is $n \in [0, N], n \in \mathbb{N}$. There are $\theta n = n_h$ high-risk agents and $(1-\theta)\, n = n_l$ low-risk agents in the pool, with θ denoting the proportion of high-risk agents, $\theta \in [0,1]$, $n_h, n_l \in \mathbb{N}$. An agent of type k enters the pool in exchange for an actuarially calculated premium π_k,

$$\pi_k = p_k L c\,, \tag{1}$$

where c is a risk loading on top of the pure premium, which enables the accumulation of risk-bearing capital. In the literature, the risk loading is often represented as an additive loading rather than a proportional one, because of the homogeneity assumption. It is then derived from the application of the law of large numbers and the central limit theorem. A multiplicative loading can also be calculated in the same manner: one can derive the loading for an average expected loss and adjust it to be proportional to the individual expected loss.

We assume that all the capital is provided by policyholders. The quantity of premiums required to cover one loss is equal to $\hat{n} = \{n \mid \theta n \pi_h + (1-\theta) n \pi_l = L\}$. Then, the minimum number of policyholders required to cover one loss is

$$\overline{n} = \min n \in \mathbb{N} \mid \theta n \pi_h + (1-\theta) n \pi_l \geq L\,, \tag{2}$$

In other words, if $\hat{n}$ is the premium quantity required to cover exactly one loss, then $\overline{n}$ is the number of policyholders required to provide those premiums: $\overline{n} = \lceil \hat{n} \rceil$. It is therefore the minimum pool size required to achieve the capacity to cover one loss.

The minimum pool size $\overline{n}$ required to cover one loss depends on the pool composition: multiple combinations of high and low risks provide the necessary amount of the risk-bearing capital. For a given number of high-risk policyholders $\overline{n}_h$, the minimum pool size $\overline{n}$ is given by

$$\overline{n} = \left\lceil \frac{L - \overline{n}_h \left(\pi_h - \pi_l\right)}{\pi_l} \right\rceil . \tag{3}$$

The premium quantity required to cover exactly one loss is $\hat{n} = \overline{n}_h + \hat{n}_l$, with $\hat{n}_l$ such that $\overline{n}_h \pi_h + \hat{n}_l \pi_l = L$. Furthermore, since $\overline{n}$ denote the minimum required number of policyholders, we must have $\overline{n}_h, \overline{n}_l \in \mathbb{N}$.

A low-risk agent has a lower loss probability than a high-risk agent. Yet, she also contributes less to the common reserves. A decrease in risk-bearing capital generated by a diminished proportion of high-risk policyholders θ cannot be balanced by an increased proportion of low-risk policyholders $(1-\theta)$, as it appears in the expression for collected premiums: $\theta n \pi_h + (1-\theta) n \pi_l$. Furthermore, insurance funds, represented by premiums here, serve their purpose when losses can actually be covered. Since it requires a precise number and combination of policyholders to cover one loss, an additional policyholder does not necessarily improve the coverage capacity.

Proposition 1. *For any given number of high-risk policyholders $\overline{n}_h$, the minimum pool size $\overline{n}$ gives the combination necessary to cover one loss at most: $\overline{n} = \left\lceil \frac{L-\overline{n}_h(\pi_h-\pi_l)}{\pi_l} \right\rceil$. Consequently,*

i) there is no direct substitution between two risk types;

ii) an additional policyholder does not always improve the risk-bearing capacity of the pool;

Consider a simple numerical illustration. Each agent faces a risk of losing one monetary unit: $L = 1$. A high-risk agent's loss probability is τ times higher than the one of a low-risk agent: $p_h = \tau p_l$, $\tau > 1$. Since an insurance premium is proportional to the expected loss, a high-risk premium is also τ times higher: $\pi_l = 0.04$ and $\pi_h = 0.08$ $(\tau = 2)$. In an homogeneous pool, the amount of premiums required to cover exactly one loss is equal to $\frac{L}{\pi_k} = \frac{1}{p_k c}$. Since $p_h = \tau p_l$, the quantity of premiums required to cover exactly one loss in a high-risk pool is τ times lower. In other words, at least τ times less policyholders are required to cover one loss in a high-risk pool. The required minimum size would be either thirteen high-risk policyholders, or twenty-five low-risk policyholders: since high-risk agents contribute twice as much as low-risk agents, it takes at least twice as much low-risk policyholders to cover the same loss. If both types are allowed in the pool, then the minimum size required to absorb one loss is given by

$$\overline{n} = \left\lceil \frac{1 - \overline{n}_h(0.08 - 0.04)}{0.04} \right\rceil . \tag{4}$$

Suppose that we already have nine high-risk policyholders in the pool. Then, the minimum pool size is sixteen, and it would be necessary to add seven low-risk policyholders in order to accumulate enough funds to cover one potential loss. A high-risk policyholder could not be replaced by a low-risk one. The expected global loss in the pool decreases when a high-risk policyholder is replaced by a low-risk one, but it might also deteriorate the risk-bearing capacity by reducing the available funds. Next, if only six low-risk agents enter the pool, one accident would still be enough to make the pool go bankrupt. In other words, the risk-bearing capacity does not improve progressively with size, because the risk-bearing capital provided by policyholders improves the risk-bearing capacity only when those funds enable the coverage of one extra loss.

Overall, the accumulation of the risk-bearing capital is related to the pool size and composition. Multiple combinations of policyholders allow to cover a given number of losses. The link between the capacity to cover losses and the global expected loss in the pool determines insurer's risk of ruin. An increase in the risk-bearing capital does not imply an improvement of the risk-bearing capacity. Thus, an additional policyholder does not systematically decrease insurer's ruin probability: the risk of ruin is not decreasing continuously with size. As we show it in the next section, the ruin probability decreases only when the capacity to cover losses is reinforced. We explicit the link between the risk of ruin and the risk-bearing capacity, which depends on the pool's composition.

3 Risk-bearing capacity and ruin probability

In this section, we examine the evolution of the insurer's ruin probability conditional on the pool composition. Further, we provide an example illustrating high-risk policyholders participation in the decrease of the risk of ruin.

Consider a simple one period model of insurer's risk. The insurer has no initial capital. He collects premiums from n policyholders at the beginning and uses these funds to cover the losses at the end of the period. The aggregate loss at the end of the period is the sum of realizations x_i of the random variable X_i , $i = 1, ..., n$ for each insured i, which takes value 0 or L. Insurer's resulting surplus S ,

$$S = \sum_{i=1}^{n} \pi_i - \sum_{i=1}^{n} x_i \, , \tag{5}$$

can be positive or negative.

In this model, insurer's risk is a probability of a negative result, $\Pr(S < 0)$, which leads to the ruin. Thus, the risk of ruin is a probability that the aggregate loss be higher than the available funds,

$$\begin{aligned} \Pr(ruin) &= \Pr\left(\sum_{i=1}^{n} x_i > \sum_{i=1}^{n} \pi_i\right) \\ &= \Pr\Big(\sum_{i=1}^{n} x_i > \theta n \pi_h + (1-\theta) n \pi_l\Big). \end{aligned} \tag{6}$$

We further illustrate the fact that this probability is non-monotonous and discontinuous. In particular, we want to emphasize the fact that the ruin probability is locally increasing with the number of policyholders, even though it ultimately decreases at the global level when the pool size grows larger.

As it was highlighted in the previous section, those characteristics of the ruin probability function are due to the capital accumulation that reinforces the risk-bearing capacity. An additional policyholder does not necessarily improve the capacity to cover losses. Hence, an increase in size generates an increase in the ruin probability, except when the risk-bearing capacity expands sufficiently to cover an additional loss. At this point, the expansion of the risk-bearing capacity produces a drop in the ruin probability. Thus, the ruin probability function has multiple discontinuity points. It is increasing on each consecutive interval, and each of them ends with a momentary decrease. Those intervals are determined by the risk-bearing capacity, which depends on the pool size and composition. Hence, each interval can be represented by the minimum number of policyholders required to cover one loss, $\overline{n}$. An upper bound of m-th interval is therefore a multiple of this number, $m\overline{n}$, which represents the number of policyholders allowing the coverage of m losses.

Proposition 2. *An increase of the pool size that permits an improvement of the loss coverage induces a one-time decrease of the ruin probability. At this point, the ruin probability is lower than on the previous intervals, except the first one:*

$\forall n = m\overline{n}$, $m \in \mathbb{N}$, $m > 1$, $j \in [1, (m-1)\overline{n}]$,

$$\Pr\left(\sum_{i=1}^{n} x_i > \theta n\pi_h + (1-\theta)n\pi_l\right) < \Pr\left(\sum_{i=1}^{n-j} x_i > \theta(n-j)\pi_h + (1-\theta)(n-j)\pi_l\right)$$

Tables 1, 2 and 3 below sum up the evolution of the ruin probability in three pools that differ in composition. Each policyholder may lose one monetary unit ($L = 1$). The premiums are the same as in the numerical example from the previous section: $\pi_l = 0.04$, $\pi_h = 0.08, c = 4, p_l = 0.01, p_h = 0.02$. The value of the risk loading is fixed arbitrarily and does not decrease with the size, so that the increase in the number of policyholders induces a decrease in the ruin probability. Another possible case is to fix the ruin probability and decrease the premiums. This point is discussed in Section 4.

Table 1: Ruin probability in the low-risk pool

Number of policyholders (n)	Total resources ($n\pi_l$)	Coverage capacity (mL)	Ruin probability
1	0.04	0	0.0100
2	0.08	0	0.0199
24	0.96	0	0.2143
25	1.00	**1**	**0.0258**
26	1.04	1	0.0277
49	1.96	1	0.0864
50	2.00	**2**	**0.0138**

Table 2: Ruin probability in the high-risk pool

Number of policyholders (n)	Total resources ($n\pi_h$)	Coverage capacity (mL)	Ruin probability
1	0.08	0	0.0200
2	0.16	0	0.0396
12	0.96	0	0.2153
13	1.04	**1**	**0.0270**
14	1.12	1	0.0310
24	1.92	1	0.0826
25	2.00	**2**	**0.0132**

Table 1 above depicts the ruin probability evolution relative to the changes in size in the low-risk pool. The ruin probability increases from 0.01 to 0.2143 on the interval $[1, 25]$. As it was shown previously, the minimum number of low-risk policyholders that enables a capacity to cover one loss is twenty-five agents ($\overline{n} = 25$). Once this pool size is reached, the ruin probability drops to 0.0258. There is a further increase until twenty-five additional policyholders join the pool. At this point, the pool reaches the upper bound of the next interval. The capacity to absorb two losses ($m = 2$) is associated with the size of $2\overline{n}$ policyholders and reduces the ruin probability to 0.0138.

In case of the high-risk pool (Table 2), the ruin probability increases from 0.02 to 0.2153, then drops to 0.027 when the pool reaches a size corresponding to thirteen policyholders. The following increase persists until twelve additional policyholders join the pool[4].

Table 3 below illustrates the case of an heterogeneous pool with a fifty-fifty representation. The ruin probability is increasing until nine participants of each type join the pool. At this point, it drops to 0.0292. It is increasing again on the next interval until the upper bound is reached, and it drops second time to 0.0142.

Table 3: Ruin probability in the pool with a fifty-fifty representation

Number of policyholders ($\theta = 0.5$)	Total resources ($\theta n \pi_h + (1-\theta) n \pi_l$)	Coverage capacity (mL)	Ruin probability
2	0.12	0	0.0298
4	0.24	0	0.0587
16	0.96	0	0.2150
18	1.08	**1**	**0.0292**
20	1.20	1	0.0356
22	1.32	1	0.0425
32	1.92	1	0.0829
34	2.04	**2**	**0.0142**

This numerical example provides an illustration for the following fact: in presence of the proportional risk loading, the ruin probability level is lower in the high-risk pool than it is in the low-risk pool of the same size. Thus, it highlights the importance of the high-risk policyholders for the accumulation of the risk-bearing capital and, by so, for the improvement of the insurer's risk-bearing capacity. In the heterogeneous pool, if the high-risk policyholders leave the insurance pool, it may deteriorate the coverage capacity and the pooling benefits for the remaining low-risk policyholders. Moreover, the situation can become unacceptable from the regulator's point of view.

[4]It is twelve because $m\overline{n} = \lceil m\hat{n} \rceil$.

Proposition 3. *In presence of the proportional risk loading, the ruin probability level is lower in the high-risk pool than it is in the low-risk pool of the same size.*

According to the legal requirements, insurers are bound to control their risk of ruin in order to conform to the policyholders protection terms. Those terms are currently set within the framework of Solvency II. According to the directive, insurers must be in line with the Solvency Capital Requirement, which is the amount of capital to be held in order to ensure an acceptable level of financial safety. The latter corresponds to the probability of solvency of 99.5% over the twelve months, thus limiting the chance of ruin to less than once in 200 cases[5].

Finally, the level of safety depends on the risk level and on the risk-bearing capital. Both depend on the pool size and composition. It is therefore possible to find a combination of policyholders $m\bar{n}$ such that any further increase in size maintains the ruin probability below a target level. If the target level is 0.5% (0.005), then at least 125 policyholders are required in case of an homogeneous low-risk pool (see Table 4 below), while only sixty-three policyholders are required in case of the high-risk pool (see Table 5 below).

Table 4: Solvency threshold in the low-risk pool

Number of policyholders	Total resources	Coverage capacity	Ruin probability
25	1.00	1	0.0258
75	3.00	3	0.0069
100	4.00	4	0.0034
125	5.00	5	**0.0017**
126	5.04	5	0.0018
149	5.96	5	**0.0041**

Table 5: Solvency threshold in the high-risk pool

Number of policyholders	Total resources	Coverage capacity	Ruin probability
13	1.04	1	0.0270
38	3.04	3	0.0069
50	4.00	4	0.0032
63	5.04	5	**0.0016**
64	5.12	5	0.0018
74	5.92	5	**0.0037**

[5]There is a indeed a difference between the concepts of "ruin probability" and "probability of insolvency": the notion of insolvency generally applies to a multi-period models. We are aware of this detail, but still use both interchangeably.

If the high-risk policyholders leave the insurance pool, it may deteriorate insurer's risk of ruin. In the next section, we discuss the implications of our observations.

4 Implications of the high risks crowding out

In this section, we discuss the implications of the high-risk policyholders crowding out.

Through the paper, we have made an assumption that an insurer calculates a fixed risk loading at the beginning of the period. The presence of the fixed loading in the insurance premium enables the accumulation of the capital. Thus, the risk of ruin ultimately decreases with the number of policyholders in the pool.

Another possibility for the insurer is to fix a desired level of ruin probability at the beginning. In this case, the risk loading is calculated based on this target level and the size of the risk loading ultimately decreases with the number of policyholders. As we have shown in the previous section, the ruin probability level is lower in the high-risk pool than it is in the low-risk pool of the same size, if the risk loading is proportional. Two implications are following: it requires more low-risk policyholders to achieve a given level of safety, or it requires a higher risk loading.

If an insurer must conform to the regulation requirements, as it is the case with the Solvency II Directive, he is bound to achieve a required level of safety. Consequently, he can either create a sufficiently big low-risk pool, or he can increase the size of the risk loading. If it is not possible to gather the required number of participants, it will be compensated by the premium increase.

This implication is particularly important for the new insurance firms that are sensible to the changes in the number of clients. For instance, the insuretech startup companies find it difficult to gather enough policyholders. Consequently, they are often affiliated to the big insurers with an established portfolio. It is also important for the insurance sectors with a limited number of potential policyholders, as it is the case for the insurance against specific risks, or artisans liability insurance. Indeed, if an insurance pool is already big, i.e., below the requirements concerning the ruin probability, the insurer can accept a decrease in his portfolio size, which may not be the case if his portfolio is small.

High risks are not bad: there are no bad risks if they are correctly estimated and priced, i.e., if each policyholder's contribution to the reserves corresponds to her individual risk level. A high risk is a bad risk if it is impossible to know whether she is one. The "lemon market" issue arises because of the hidden knowledge[6].

[6] Akerlof (1970).

For instance, Powers (2006) observes an apparent lack of proper benefits from the risk pooling in companies with a big underwriting volume. He points out that the premium-surplus ratio does not properly increase with the number of policyholders, as it should have be the case if the law of large numbers was working fine. He argues that those findings might be due to the negative effects accompanying an increase in size, in particular the possibility of an increased underwriting error and a "reduced quality of risks".

The "lemons" and the underwriting errors should be less of an issue if each individual risk is assessed correctly. If the technological advance and the access to the larger volumes of information may allow to determine the individual risk level of each policyholder, the conclusion should go in the opposite direction than the one encountered in the media: the high risk should not be dismissed.

A possible research question related to the use of information on the individual risk level is its potential to provide incentives to invest in loss prevention. Self-protection can be enhanced by creating explicit recommendations relative to the prevention effort. It can therefore allow to decrease the individual risk level through the changes in risk factors that are endogenous.

5 Conclusion

In this paper, we show that multiple combinations of policyholders allow the insurer to achieve a given level of risk-bearing capacity, and so a given level of safety. We highlight the fact that an additional policyholder increases the funds, but not necessarily the capacity to cover the losses. Consequently, the risk of ruin does not necessarily decrease when an additional policyholder joins the pool. The ruin probability decreases only when the coverage capacity is expanded, and each decrease is associated with a combination of policyholders determined by their risk type.

High-risk policyholders' contribution to the risk-bearing capital is higher than the contribution of the low-risk policyholders. Consequently, other things being equal, a higher number of low-risk policyholders is required in order to achieve a given level of safety. Legal requirements on the acceptable ruin probability level are particularly important in the case of the new insurance firms and the small firms operating in the sectors with a limited number of potential policyholders.

In theory, a difficulty to create a sufficiently big insurance pool can be compensated by an increase in premiums for the existing policyholders, which is not always an option in the competitive setting. If the competition is high, the insurer cannot increase the premiums without losing the policyholders. Consequently, insurers often merge together in order to meet the solvency constraints. It might be optimal to create a mix of policyholders in order to achieve the necessary level of the ruin probability.

References

Akerlof, George A. (1970). “The Market for "Lemons": Quality uncertainty and the Market Mechanism”. In: *The Quarterly Journal of Economics* 84(3), pp. 488–500.

Borch, Karl (1962). “Equilibrium in a Reinsurance Market”. In: *Econometrica* 30(3), pp. 424–444.

Cummins, J David (1974). “Insurer’s Risk: A Restatement”. In: *The Journal of Risk and Insurance* 41(1), pp. 147–157.

Cummins, J David (1991). “Statistical and financial models of insurance pricing and the insurance firm”. In: *The Journal of Risk and Insurance* 58(2), pp. 261–302.

Denuit, Michel M., Louis Eeckhoudt, and Mario Menegatti (2011). “A Note on Subadditivity of Zero-Utility Premiums”. In: *ASTIN Bulletin: The Journal of the IAA* 41(1), pp. 239–250.

Feller, William (1968). *An Introduction to Probability Theory and Its Applications.* John Wiley & Sons.

Gatzert, Nadine and Hato Schmeiser (2012). “The merits of pooling claims revisited”. In: *The Journal of Risk Finance* 13(3), pp. 184–198.

Houston, David B. (1964). “Risk, insurance, and sampling”. In: *The Journal of Risk and Insurance* 31(4), pp. 511–538.

Klein, Florian and Hato Schmeiser (2018). “Optimal Risk Pools under Changing Volatility and Heterogeneous Risk Classes”. In: Working Paper on Risk Management and Insurance.

Powers, Michael R. (2006). “An insurance paradox”. In: *The Journal of Risk Finance* 7(2), pp. 113–116.

Powers, Michael R., Emilio C. Venezian, and Iana B. Jucá (2003). “Of happy and hapless regulators: the asymptotics of ruin”. In: *Insurance: Mathematics and Economics* 32(2), pp. 317–330.

Smith, Michael L. and Stephen A. Kane (1994). “The Law of Large Numbers and the Strength of Insurance”. In: *Huebner International Series on Risk, Insurance and Economic Security* 18.

Venezian, Emilio C. (1984). “Efficiency and equity in insurance”. In: *The Journal of Risk and Insurance* 51(2), pp. 190–204.