Decisions in Economics and Finance
https://doi.org/10.1007/s10203-026-00569-6

Shapley risk sharing in peer-to-peer insurance

Gian Paolo Clemente[1] · Susanna Levantesi[2] · Gabriella Piscopo[3]

Received: 15 October 2025 / Accepted: 18 March 2026

Abstract
Peer-to-peer (P2P) insurance is an innovative model that leverages digital technology to connect individuals with similar insurance needs, forming a pool to share risks. This paper introduces a P2P insurance framework in which participants pay an ex-ante contribution determined by the Shapley value, where Value-at-Risk of the difference between the expected and realized total loss of the network (i.e., the Profit and Loss) is used as the risk measure. Under standard assumptions commonly used in non-life insurance for the aggregate claim amount, we derive closed-form expressions for the Shapley value. The model includes a cashback mechanism to ensure that all participants contribute equally to covering realized losses. We demonstrate the practical implementation of the model by applying it to a portfolio of motor other damage insurance policies.

Keywords Peer-to-peer insurance · Shapley value · Risk measure · Non-life insurance

1 Introduction

The rapid advancement of digital technologies, ranging from big data, the Internet of Things, and digital platforms to blockchain and cloud computing, has profoundly reshaped the insurance sector, catalyzing process automation and the emergence of new business models (EIOPA 2020, 2023). Within this context, peer-to-peer (P2P)

✉ Gian Paolo Clemente
gianpaolo.clemente@unicatt.it

Susanna Levantesi
susanna.levantesi@uniroma1.it

Gabriella Piscopo
gabriella.piscopo@unina.it

1 Department of Mathematics for Economics, Finance and Actuarial Science, Università Cattolica del Sacro Cuore, Milano, Italy

2 Department of Statistics, Sapienza University of Rome, Rome, Italy

3 Department of Economics and Statistical Science, University of Naples Federico II, Naples, Italy

Published online: 19 April 2026

insurance has gained increasing scholarly, regulatory, and industrial attention as one of the most promising applications of InsurTech.

From a regulatory standpoint, the *European Insurance and Occupational Pensions Authority* (EIOPA) has played a pivotal role. In its 2019 report, the authority examined whether P2P risk-sharing qualifies as an insurance activity and should therefore comply with sector-specific regulation, recommending the application of the principle of proportionality. This principle recognizes P2P risk-sharing as an insurance activity depending on the scope of implementation and the role of intermediaries (EIOPA 2019). EIOPA has since broadened the discussion in its digitalization strategy and subsequent reports (EIOPA 2020, 2023), situating P2P initiatives within the wider digital transformation of European insurance markets. In the United States, where insurance regulation is primarily state-based, the *National Association of Insurance Commissioners* (NAIC) has significantly shaped the debate on P2P insurance. Through its Innovation, Cybersecurity, and Technology Committee, the NAIC has highlighted both opportunities (such as enhanced transparency and reduced costs) and risks (including solvency, governance, and compliance with state law). Publications from the NAIC, including issue briefs and the Center for Insurance Policy and Research topic on P2P insurance, provide regulatory guidance on how to balance innovation with consumer safeguards in a fragmented environment (NAIC 2021). A recurrent issue in the literature on P2P insurance concerns the degree of uncertainty faced by participants and the limited scope of regulatory protection compared to traditional insurance products. Many P2P schemes operate outside the perimeter of standard insurance regulation or fall into regulatory grey areas, with important implications for solvency protection, governance, and consumer rights (EIOPA (2019); Ostrowska and Ziemiak (2020); NAIC (2021)). In particular, participants are often not covered by guarantee funds or policyholder protection schemes, and capital requirements may be absent or significantly lower than those imposed on licensed insurers. While digital platforms and smart contracts can enhance transparency and reduce transaction costs, they do not fully compensate for the lack of prudential oversight and may increase exposure to systemic and tail risks under adverse scenarios (Sun et al. (2020); Denuit and Robert (2021a)). Recent regulatory reports emphasize the need for proportionate regulation that balances innovation with consumer protection, especially as P2P models scale up and increasingly resemble traditional insurance arrangements ((EIOPA 2020, 2023); Global InsurTech Report (2023)). The regulatory debate has followed the development of industry practice and different kinds of P2P insurance across many jurisdictions have been proposed. In Europe, *Friendsurance* in Germany pioneered the broker-model structure in the early 2010s, enabling small groups of policyholders to share a portion of their premiums to cover minor claims, with unused funds redistributed as cashback. Another example is *InsurePal*, a Swiss decentralized platform founded in 2015, which leveraged blockchain technology and reputation mechanisms to foster trust and risk-sharing (Sun et al. 2020). Although innovative, these initiatives faced issues of scalability and regulatory scrutiny. In Asia, the Chinese platform *Xiang Hu Bao* attracted over 100 million participants within two years by pooling small contributions to cover critical illness costs. Structured as a mutual-aid health platform rather than a licensed insurance product, it exemplified the potential of digital ecosystems in scaling P2P schemes but also highlighted the blurred boundaries between mutual aid and insurance, eventually prompting stricter oversight by regulators. In the United States, P2P initiatives have been more limited due to the state-based regulatory framework. *Lemonade*, founded in 2015, initially presented itself as a P2P insurer but later transitioned into a fully licensed carrier while retaining some P2P-inspired features, such as the redistribution of unclaimed premiums to charities. This evolution illustrates the difficulty of sustaining a pure P2P model in the U.S. context.

Against these experiences, the classification of P2P arrangements remains a key point of debate. Three main models are commonly distinguished: the *self-governing model*, where a community of peers autonomously manages all insurance functions (Denuit and Robert 2021b); the *broker model*, the most widespread structure, in which a broker manages the group via a digital platform, participants pay an entry fee split between a common fund and an insurer, losses up to a threshold are covered by the group while larger ones are reinsured, and surpluses are refunded as cashback while deficits are borne by the insurer (Rego and Carvalho 2019); and the *carrier model*, an extension of the broker model that effectively operates as a fully licensed insurance company entirely managed on digital platforms.

The academic literature on these models has expanded considerably in recent years. In the case of self-governing arrangements, Denuit and Dhaene (2012) introduced the Conditional Mean Risk (CMR) allocation rule, later refined by Denuit (2019) and Denuit and Robert (2020), who emphasized its Pareto-optimality under co-monotonic losses. Further contributions include extensions with ex-ante contributions and safety loading (Levantesi and Piscopo 2022), multivariate adaptations (Abdikerimova and Feng 2022), and systematic comparisons of allocation rules (Denuit et al. 2022). Building on this tradition, Feng et al. (2023) proposed an optimal allocation mechanism consistent with Pareto efficiency and actuarial fairness. In the case of broker-based models, where premiums are collected upfront, Denuit and Robert (2021a) analyzed entry fees analogous to classical premium principles, while Charpentier et al. (2021) studied collaborative systems structured around social networks. Cooperative game theory has also been applied: Clemente et al. (2023) developed a Shapley Value-based cashback distribution rule, and Galeotti and Rabitti (2024) employed the Shapley Value to price reinsurance contracts within P2P frameworks. Our work is related to Galeotti and Rabitti (2024), who study Shapley value allocations of tail risk measures from a theoretical perspective based on majorization arguments. While we also rely on the Shapley value, our contribution differs in several important aspects. We embed the allocation mechanism within a broker-based P2P insurance framework, provide an institutional interpretation of the allocated tail risk as a safety loading added to expected losses, and derive explicit contribution rules under standard actuarial loss models. The two approaches are therefore complementary. While Galeotti and Rabitti (2024) offer a general theoretical foundation for Shapley allocations of tail risk, the present paper emphasizes economic interpretation, institutional design, and applied actuarial relevance.

Despite the diversity of approaches, a recurring theme emerges: P2P insurance revives the principle of mutuality in a digital environment. Yet, as noted by Ostrowska and Ziemiak (2020), its regulatory status remains uncertain in many jurisdictions, and the absence of a universal definition raises concerns about governance and consumer protection. Emerging technologies facilitate transparency and efficiency but simultaneously complicate oversight, thereby underscoring the need for proportionate regulation.

This paper contributes to the ongoing debate by extending the cooperative game framework to broker-based P2P insurance. Specifically, we introduce a safety loading into the entry premium calculation and propose a Shapley value-based rule for cash-back distribution, thereby linking actuarial valuations to practical market mechanisms. We develop a risk-sharing mechanism where the total expected loss is equally shared, while deviations from expectations (uncertainty) are allocated using Shapley values on a Value-at-Risk basis.

While several alternative risk measures widely used in actuarial science, such as distortion risk measures (Wang 2000), spectral risk measures (Acerbi 2002), and more general convex and utility-based risk measures (Bellini et al. 2014; Föllmer and Schied 2016), offer appealing theoretical properties, their integration into a Shapley-value allocation framework rarely leads to closed-form expressions for marginal contributions, except under restrictive distributional assumptions. Because these measures are typically defined as weighted integrals of quantiles or other non-linear functionals of loss distributions, the resulting marginal contributions generally require numerical or simulation-based methods. By contrast, under standard non-life insurance assumptions, Value-at-Risk retains analytical tractability and admits explicit expressions for both aggregate risk and marginal contributions, which motivates its adoption in the present paper.

Grounded in cooperative game theory and probabilistic modeling, our approach balances simplicity (equal sharing) and risk-based allocation, thus advancing the literature on risk-sharing and capital contribution mechanisms in P2P insurance. By leveraging the Shapley value, our framework captures each participant's marginal contribution to risk, ensures fairness by satisfying the axioms of efficiency, symmetry, dummy, and additivity, and provides a transparent and justifiable allocation of risk-derived capital requirements or contributions.

The remainder of the paper is organized as follows. Section 2 is devoted to the description of the application of the Shapley value in P2P insurance context. In Section 3 the theoretical framework is outlined, while in Section 4 numerical results are presented and discussed. Final remarks are offered in Section 5.

2 On the Shapley value in P2P insurance

The Shapley value (Shapley 1953) is a concept from cooperative game theory used to fairly allocate the total output of a game among its participants. Let N denote the set of n players in a cooperative game, and let S with $S \subseteq N$ represent a coalition of players, where $|S| = s$, with $s \leq n$, is the number of players in the coalition. The set $N \setminus \{i\}$ refers to all possible coalitions excluding player i. Given the characteristic function of the game $\mu(S)$, the Shapley value for player i quantifies their contribution to the total outcome of the game, taking into account the risk they confer to the entire set of players N:

$$\varphi_i(\mu) = \sum_{S \subseteq N \setminus \{i\}} \frac{s!(n-s-1)!}{n!} \left[\mu(S \cup \{i\}) - \mu(S)\right]. \tag{1}$$

where $[\mu(S \cup \{i\}) - \mu(S)]$ represents the marginal contribution of player i to coalition S. This marginal contribution is averaged over all possible permutations of the players, ensuring that each coalition is weighted equally according to the number of players involved. The game-theoretic approach adopted in this paper belongs to the class of *cooperative games with transferable utility*. The P2P insurance pool is modeled as a cooperative game in which the set of players corresponds to the participants in the scheme, and the characteristic function assigns to each coalition a measure of aggregate risk or cost associated with the losses of its members. Game theory is therefore not used to model strategic interaction or equilibrium behavior, but rather as a normative tool to address a risk and cost allocation problem. In particular, we employ the Shapley value as an allocation rule to decompose the aggregate risk of the pool into individual contributions. The Shapley value allocates risk according to each participant's marginal contribution across all possible coalitions and satisfies key axiomatic properties, efficiency, symmetry, additivity, and the dummy player property, which are desirable in insurance and capital allocation contexts. This application of cooperative game theory is well established in actuarial science and financial risk management, especially in the context of capital and risk allocation. In a cooperative game, players contribute in different ways to the achievement of a final outcome but collaborate to achieve a payoff that is at least as large as what they could have obtained by acting independently. The Shapley value provides a fair method for distributing the total gain among the players (Roth 1988). Shapley (1953) demonstrated that the rule in Eq. 1 is the unique optimal risk allocation that satisfies the axiomatic properties of efficiency, symmetry, additivity, and the axiom of dummy players.

In the context of P2P insurance, small groups of individuals agree to share risks among themselves, potentially bypassing traditional insurers or supplementing conventional insurance coverage. Here, the Shapley value offers a rigorous and equitable method for allocating risk or costs based on each participant's marginal contribution to the collective risk. Unlike simple proportional allocation rules, the Shapley value accounts not only for the expected size of a participant's loss but also for how their individual risk contributes to the overall risk of the pool. This ensures that each member's premium or contribution is aligned with their actual exposure to risk. For example, a participant whose losses are highly correlated with the total loss (due to factors such as geography, behavior, or other influences) will bear a larger share of the pool's safety loading, as determined by their marginal contribution to the total risk. Although individual contributions in our framework ultimately depend on the aggregate loss of the network, and hence on the grand coalition N, the explicit consideration of all coalitions $S \subseteq N$ remains essential from a methodological perspective. In line with standard cooperative game theory, the Shapley value is defined as the average of marginal contributions taken over all possible coalitions, which serve as hypothetical constructs rather than economically implementable groupings. In the present context, sub-coalitions do not represent alternative risk-sharing arrangements or feasible deviations from the pool; instead, they serve as a formal device for assessing how each participant contributes to aggregate risk through diversification and dependence effects. Consequently, while the grand coalition N is the only economically relevant coalition, since both total losses and total contributions are defined at the aggregate level, the full coalition structure is required to derive an allocation rule that satisfies the axioms of efficiency, symmetry, additivity, and the dummy player property. The Shapley value thus provides a principled means of decomposing the aggregate risk of the pool into individual contributions, even when the resulting contribution rule can be expressed solely as a function of the aggregate loss.

By linking contributions to marginal risk, the Shapley value approach helps mitigate adverse selection. Members with higher risk correlations are required to pay more, while risk-reducing behavior is incentivized. For instance, diversifying or reducing exposure can lower one's marginal contribution, and consequently, their premium. A transparent, risk-based approach is particularly important in P2P insurance schemes, where trust among participants is key. The Shapley value, rooted in well-established cooperative game theory, provides a theoretically sound and transparent rule that can enhance confidence in the system. Specifically, in the context of P2P insurance, it plays a central role in risk sharing by decomposing the aggregate risk of the pool into individual contributions based on marginal impact. Rather than allocating losses or premiums proportionally to expected losses alone, the Shapley value accounts for how each participant affects the total risk through volatility and dependence structures. In this way, diversification benefits are endogenously reflected in the allocation: participants who reduce aggregate risk bear a smaller share of the risk load, while those who increase tail risk contribute more. The Shapley value thus provides a transparent and economically meaningful mechanism for sharing risk among peers. The Shapley value has been applied in P2P insurance in different ways. In *Premium allocation*, it has been used to ensure a fair distribution of the collective premium needed to cover expected losses (see (Clemente et al. 2024) and Galeotti and Rabitti (2024) for pricing peer-to-peer reinsurance contracts). Alternatively, it has been used to assess the *Residual amount distribution*, equitably distributing any surplus or deficit (after claims are paid), in proportion to each member's marginal contribution to total risk (Clemente et al. 2023).

3 The theoretical framework

The individual and total loss. Let n denote the number of participants in the P2P scheme. Each participant i (for $i = 1, \ldots, n$) faces a loss represented by a non-negative random variable L_i. The total loss of the group, denoted by X, is defined as the sum of individual losses:

$$X = \sum_{i=1}^{n} L_i \tag{2}$$

The Profit and Loss (P&L). We define the random variable Profit and Loss (P&L), denoted by Y, as the difference between the expected total loss of the network $\mathbb{E}(X)$ and the random total loss X:

$$Y = \mathbb{E}(X) - X \quad \text{so that} \quad \mathbb{E}(Y) = 0. \tag{3}$$

Using Eq. (2), we rewrite Eq. (3) as:

$$Y = \mathbb{E}\left(\sum_{i=1}^{n} L_i\right) - \sum_{i=1}^{n} L_i$$

which simplifies to

$$Y = \sum_{i=1}^{n} \left[\mathbb{E}(L_i) - L_i\right]. \tag{4}$$

We can express the total P&L as

$$Y = \sum_{i=1}^{n} Y_i. \tag{5}$$

where Y_i is the r.v. P&L of the generic policyholder i.

We assume that Y follows a Gaussian distribution with mean $\mu = 0$ and standard deviation σ, i.e., $Y \sim \mathcal{N}(0, \sigma^2)$. Given Eq. 3, this implies that the total loss X also follows a Gaussian distribution with mean $\mathbb{E}[X]$ and variance σ^2. This assumption of normality for the P&L variable can be theoretically justified by the Central Limit Theorem and is frequently employed in actuarial science. For example, Klugman et al. (2019) discuss the practical utility of the normal approximation for aggregate loss distributions, especially for large portfolios. Hardy (2006) presents this approximation as a classical method for estimating aggregate claim distributions, while also highlighting its limitations in cases involving heavy-tailed claim sizes or small portfolios. Beyond this standard assumption, Section 3.1 addresses scenarios where the distribution of Y is skewed.

3.1 The contribution rule

Participants agree on a fair and transparent rule that determines how much each participant must contribute, denoted by a positive monetary amount $h_i(X) > 0$ to cover the scheme's total loss.

In a purely self-organized P2P system, members do not pay any upfront fees; instead, they contribute *ex-post* to the realized aggregate loss according to a predefined cooperative risk-sharing rule. Conversely, in a broker-based model, participants' contributions are typically paid *ex-ante*.

Individual contribution rules from the literature. Several risk-sharing rules have been proposed in the literature to define both ex-ante contributions (typical of broker-based models) and ex-post contributions (common in self-organized schemes). Examples include:

- *Ex-post contributions.* Denuit and Dhaene (2012) and Denuit (2019) introduced the *Conditional Mean Rule* (CMR), under which each participant i contributes an amount equal to the expected value of their individual loss, conditional on the realized total loss $X = x$:

  $$h_i^{\text{CMR}}(x) = \mathbb{E}(L_i \mid X = x).$$

  This rule ensures that the individual contributions exactly sum to the total realized loss x:

  $$\sum_i h_i^{\text{CMR}}(x) = \sum_i \mathbb{E}(L_i \mid X = x) = x,$$

  leaving neither surplus nor deficit.
- *Ex-ante contributions.* Clemente et al. (2023) proposed a rule based on the unconditional expectation of individual losses:

  $$h_i(X) = \mathbb{E}(L_i).$$

  In this case, the total contributions satisfy $\sum_{i=1}^n h_i(X) = \mathbb{E}(X)$, ensuring that total contributions match the expected total loss.
- *Ex-ante contributions with safety loading.* Clemente et al. (2024) extended this approach by incorporating a safety loading term based on the Shapley value to reflect participants' marginal contributions to total risk, measured by the total variance of losses:

  $$h_i(X) = \mathbb{E}(L_i) + \alpha \cdot \varphi(\text{var}(X))$$

  where $\varphi(\text{var}(X)) = \text{cov}(L_i, X)$. Here, total contributions become $\sum_{i=1}^n h_i(X) = \sum_{i=1}^n \mathbb{E}(X) + \alpha \cdot \text{var}(X)$, thus incorporating a risk premium proportional to the variability of total losses.

The proposed individual contribution rule. Considering a broker-based P2P scheme, we propose an individual contribution rule that integrates both expected losses and the deviations of realized losses from expectations within a risk-sharing framework. This rule is sensitive not only to the expected loss but also to each participant's risk profile, measured via the Shapley value. The individual contribution $h_i(x)$ is defined as:

$$h_i(X) = \frac{\mathbb{E}(X)}{n} + \varphi_i\big(VaR_\alpha(Y)\big) \tag{6}$$

where $\mathbb{E}(X) = \sum_{i=1}^n \mathbb{E}(L_i)$, and $\varphi_i(VaR_\alpha(Y))$ denotes the Shapley value allocated to participant i, based on the risk measure $VaR_\alpha(Y)$, which is the Value-at-Risk (VaR) of P&L at a confidence level α. This formulation combines an equal sharing of the expected total loss, $\frac{\mathbb{E}(X)}{n}$, ensuring that all participants bear the baseline cost, with a fair allocation of the risk arising from deviations. The latter distributes $VaR_\alpha(Y)$ among participants according to cooperative game theory principles, providing a game-theoretically justified risk-sharing scheme. Under the Gaussian assumption, $VaR_\alpha(Y)$ is linear in the standard deviation and can be expressed as:

$$VaR_\alpha(Y) = -\Phi^{-1}(\alpha) \cdot \sigma(Y)$$

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. Therefore, we obtain:

$$\varphi_i(VaR_\alpha(Y)) = -\Phi^{-1}(\alpha) \cdot \varphi_i(\sigma(Y))$$

By exploiting the linearity and additivity properties of the Shapley value, and using the approximation proposed by Hagan et al. (2023):

$$\varphi_i(\sigma(Y)) \approx \rho(Y_i, Y)\sigma(Y_i),$$

we obtain the following approximation for the individual Shapley-based contribution to the VaR:

$$\varphi_i(VaR_\alpha(Y)) \approx -\Phi^{-1}(\alpha)\rho(Y_i, Y)\sigma(Y_i), \tag{7}$$

where $\rho(Y_i, Y) = \frac{\text{cov}(Y_i, Y)}{\sigma(Y_i)\sigma(Y)}$ is the Pearson correlation coefficient between participant i's deviation Y_i and the total P&L Y. Equation (7) attributes to each participant a share of the total risk based on both their expected deviation and their correlation with the overall risk of the network. Substituting this back into Eq. (6) yields a simplified individual contribution rule:

$$h_i(X) = \frac{\mathbb{E}(X)}{n} - \Phi^{-1}(\alpha)\,\rho(Y_i, Y)\,\sigma(Y_i). \tag{8}$$

Equation (8) defines the individual contribution $h_i(X)$ of participant i to the total loss-sharing scheme. This expression consists of two components:

- *Baseline contribution*: the term $\frac{\mathbb{E}(X)}{n}$ represents an equal division of the expected total loss among all n participants, serving as a fairness benchmark under symmetric expectations.
- *Risk adjustment term*: the second term adjusts each participant's share according to their individual risk profile:
  - $\Phi^{-1}(\alpha)$ is the quantile of the standard normal distribution at confidence level α, representing the VaR threshold.
  - $\rho(Y_i, Y)$ is the Pearson correlation coefficient between the individual deviation Y_i and the total deviation Y.
  - $\sigma(Y_i)$ denotes the standard deviation of Y_i, reflecting the variability in participant i's deviation from their expected loss.

This allocation balances fairness and risk sensitivity. Participants whose losses are more positively correlated with the total loss (i.e., higher $\rho(Y_i, Y)$) are assigned higher contributions, while those with lower or negative correlation benefit from diversification and pay less. The rule thus rewards participants who reduce overall risk and penalizes those who increase it.

Noting that $\sigma(Y_i) = \sigma(L_i)$, and expressing the correlation coefficient $\rho(Y_i, Y)$ as:

$$\rho(Y_i, Y) = \frac{\sum_{j=1}^{n} \operatorname{cov}(L_i, L_j)}{\sigma(L_i)\sigma(X)}$$

we rewrite Eq. (8) as:

$$h_i(X) = \frac{\mathbb{E}(X)}{n} - \Phi^{-1}(\alpha) \frac{\sum_{j=1}^{n} \operatorname{cov}(L_i, L_j)}{\sigma(X)}. \tag{9}$$

In the special case where all individual losses are uncorrelated, the correlation reduces to $\rho(Y_i, Y) = \frac{\sigma(L_i)}{\sigma(X)}$, and Eq. (8) simplifies to: $h_i(X) = \frac{\mathbb{E}(X)}{n} - \Phi^{-1}(\alpha) \frac{\operatorname{var}(L_i)}{\sigma(X)}$.

When participants' loss L_i are not identical distributed, the term $\frac{E(X)}{n}$ in formulas (6)–(9) should not be interpreted as an actuarially fair premium in the classical stand-alone sense, that is, as a contribution proportional to $\mathbb{E}[L_i]$. Rather, it represents the mutualistic component of the contribution rule, reflecting the cooperative nature of P2P insurance schemes. Indeed, in P2P arrangements, participants typically agree *ex-ante* to pool expected costs at the group level, independently of individual risk characteristics. This interpretation is consistent with the mutual insurance principle and with a broad strand of the P2P literature, where baseline contributions are decoupled from individual heterogeneity and fairness is instead enforced through the allocation of deviations from expectations or residual risk (Denuit and Dhaene 2012; Denuit 2019; Denuit et al. 2022). Similar mutualistic baselines are also adopted in broker-based P2P models, where a certain degree of redistribution is accepted in exchange for transparency and collective risk sharing (Clemente et al. (2023)).

Actuarial fairness is ensured by the second term in formula (9), which allocates a prudential loading reflecting the uncertainty surrounding the aggregate loss. This term depends explicitly on individual risk profiles through volatility and dependence with the aggregate loss, as captured by the Shapley value. Participants with higher expected losses, greater variability, or stronger correlation with the total loss are allocated a larger share of this loading, while participants who contribute more to diversification benefit from a lower adjustment. This notion of fairness, based on marginal contributions to aggregate risk rather than on stand-alone expected losses, is well established in the literature on capital and risk allocation (Hagan et al. 2023; Mango 1998; Shapley 1953) and has recently been applied to P2P insurance (Clemente et al. 2024; Galeotti and Rabitti 2024).

From this perspective, the proposed contribution rule naturally decomposes into two conceptually distinct components: a homogeneous component that sustains mutuality and solidarity within the pool, and a heterogeneous component that ensures actuarial fairness in the allocation of risk. This decomposition is particularly well suited to broker-based P2P schemes, where mutuality and redistribution coexist with the need for risk-sensitive pricing to ensure incentives, sustainability, and solvency (Denuit and Robert 2021a; Feng et al. 2023).

Cornish-Fisher approximation to the VaR. When the distribution of Y is skewed, the Cornish–Fisher expansion (Cornish and Fisher 1938) provides a useful method to approximate the quantiles of a non-normal distribution, expressing the VaR as a correction to the VaR under the normality assumption. This expansion accounts for deviations from normality by incorporating higher-order moments—specifically, skewness and kurtosis—into the quantile estimation. These moments can be estimated using the sample skewness S and sample kurtosis K, which can then be substituted to approximate the quantiles of the true distribution.

The Cornish-Fisher approximation to the VaR at confidence level α is given by:

$$\mathrm{VaR}_{\alpha,CF}(Y) = -z_{\alpha,CF} \cdot \sigma(Y) \tag{10}$$

where $z_{\alpha,CF}$ is the adjusted quantile defined as:

$$z_{\alpha,CF} = z_\alpha + \frac{1}{6}(z_\alpha^2 - 1)S + \frac{1}{24}(z_\alpha^3 - 3z_\alpha)K - \frac{1}{36}(2z_\alpha^3 - 5z_\alpha)S^2$$

with the components:

- $z_\alpha = \Phi^{-1}(\alpha)$, the α-quantile of the standard normal distribution,
- S, is the skewness of the distribution of Y,
- K is the excess kurtosis of of the distribution of Y.

This approximation reduces to the Gaussian case when $S = 0$ and $K = 0$, thus recovering the classical normal VaR formula. The Cornish-Fisher expansion provides a systematic correction that improves VaR estimates for distributions exhibiting asymmetry or heavy tails.

Consequently, when Y is not normally distributed, the Shapley value–based individual contribution adjustment (originally given in Eq. 7) can be modified to:

$$\varphi_i(VaR_\alpha(Y)) \approx -z_{\alpha,CF} \cdot \rho(Y_i, Y)\sigma(Y_i). \tag{11}$$

where $z_{\alpha,CF}$ replaces the normal quantile $-\Phi^{-1}(\alpha)$, thus incorporating skewness and kurtosis corrections into the risk allocation.

3.2 Further theoretical considerations on the ex-ante contribution

Let consider a participant with a *Constant Absolute Risk Aversion (CARA) utility function* (see, e.g., Dhaene et al. (2012) and Landsberger and Meilijson (1999) for the use of CARA in insurance):

$$u(w) = -\exp(-\gamma w), \qquad \gamma > 0,$$

where w is wealth (or negative loss) and γ is the coefficient of absolute risk aversion.

The *certainty equivalent* CE_i of a random loss L_i is defined as the deterministic amount satisfying:

$$u(-\mathrm{CE}_i) = \mathbb{E}[u(-L_i)].$$

In case L_i is normally distributed, using the moment-generating function (mgf) we have:

$$\mathbb{E}[u(-L_i)] = -\,\mathbb{E}[\exp(\gamma L_i)] = -\exp\left(\gamma \mathbb{E}[L_i] + \tfrac{1}{2}\gamma^2\sigma^2[L_i]\right).$$

Equating the two expressions:

$$-\exp(\gamma\,\mathrm{CE}_i) = -\exp\left(\gamma \mathbb{E}[L_i] + \tfrac{1}{2}\gamma^2\sigma^2[L_i]\right),$$

which implies:

$$\mathrm{CE}_i = \mathbb{E}[L_i] + \frac{\gamma}{2}\,\sigma^2[L_i],$$

that is, risk-averse participants require a higher premium to accept the same risk. It is also noteworthy that we simplified assuming that L_i is gaussian but for positively skewed losses, the certainty equivalent increases relative to the Gaussian case. Indeed, using Taylor expansion of the mgf, additional positive terms depending on skewness or higher cumulants can be inserted.

Paying the contribution is preferable to bearing own losses if the certainty equivalent of the contribution is less than the certainty equivalent of self-insurance:

$$h_i(X) \le \mathrm{CE}_i(L_i).$$

Substituting, we have:

$$\frac{\mathbb{E}(X)}{n} - \Phi^{-1}(\alpha)\,\rho(Y_i, Y)\,\sigma(Y_i) \;\le\; \mathbb{E}[L_i] + \frac{\gamma}{2}\sigma^2(L_i),$$

which can be rearranged as:

$$-\Phi^{-1}(\alpha)\,\rho(Y_i, Y)\,\sigma(L_i) \le \left(\mathbb{E}[L_i] - \frac{\mathbb{E}[X]}{n}\right) + \frac{\gamma}{2}\sigma^2(L_i).$$

This inequality ensures that the contribution under the scheme does not exceed the certainty-equivalent loss for the participant.

We can determine the confidence level α^* at which the participant is *exactly indifferent* between paying the P2P contribution and bearing their own loss by setting the inequality to equality:

$$-\Phi^{-1}(\alpha^*)\,\rho(Y_i, Y)\,\sigma(L_i) = \left(\mathbb{E}[L_i] - \frac{\mathbb{E}[X]}{n}\right) + \frac{\gamma}{2}\sigma^2(L_i)$$

Solving for α^* yields

$$\alpha^* = \Phi\left(- \frac{\left(\mathbb{E}[L_i] - \frac{\mathbb{E}[X]}{n}\right) + \frac{\gamma}{2}\sigma^2(L_i)}{\rho(Y_i, Y)\,\sigma(L_i)} \right),$$

which corresponds to a lower-tail quantile of Y_i associated with "bad scenarios" (large realized losses). Since the scheme uses a fixed value of α, defined by the broker, all the values of α such that $\alpha \geq \alpha^*$ satisfy the individual rationality condition. As the risk aversion parameter γ increases, the threshold confidence level α^* decreases, reflecting that more risk-averse participants require the scheme to cover a larger fraction of adverse outcomes to ensure individual rationality. Conversely, less risk-averse participants are willing to join at higher α values, corresponding to less conservative coverage.

Additionally, if all participants have identical expected losses and their losses are independent, we obtain

$$\alpha^* = \Phi\left(- \frac{\gamma}{2} \frac{\sigma(L_i)}{\rho(Y_i, Y)} \right) = \Phi\left(- \frac{\gamma\sigma(X)}{2} \right)$$

In this case, it is easy to verify that standard lower-tail quantiles can satisfy individual rationality.

When expected losses or variances differ, the value of α^* varies across participants. Those with higher exposure relative to the pool or higher risk require larger contributions, consistent with the risk-sharing principle. Additionally, the scheme transfers mainly the systematic component of risk (through $\rho(Y_i, Y)\sigma(L_i)$), while idiosyncratic deviations are diversified across the pool, typically reducing the effective certainty-equivalent loss compared to self-insurance.

3.3 The cashback distribution

In the proposed P2P insurance framework, unexpected losses - defined as deviations of the realized aggregate loss from its expected value - are borne entirely by the participants of the pool, with no external insurer or reinsurer assuming these risks. The role of the Shapley value applied to a VaR measure is to allocate *ex-ante* the cost of uncertainty, by determining how much each participant contributes to the risk margin associated with aggregate tail risk. Once losses are realized, any surplus or deficit relative to expected losses is redistributed among participants through the cashback mechanism. This ex-post distribution ensures that the realized aggregate loss is fully covered by the pool, while remaining consistent with the ex-ante Shapley-based allocation of risk. Hence, unexpected losses are collectively borne by participants, but their allocation is anticipated and priced ex-ante according to each participant's marginal contribution to aggregate risk. Therefore, the collected individual contributions are used to cover the total loss experienced by the group. Once the loss is realized, if any amount remains, it is distributed among the members as a cashback according to a specific rule. Conversely, if the initial contributions are not sufficient, the members are required to pay an additional sum according to the same distribution mechanism.

In this setting, let K denote the residual amount after the payment for loss, which can take both positive or negative values. If positive, it represents a cashback, while if negative, it corresponds to an additional sum to be paid to cover losses exceeding the initial contributions. Note that K is a random variable depending on the realizations of total loss and is not fixed *ex-ante* as in the classical insurance framework dealing with capital distribution. Denoting by x the realized total loss, we define the amount K as:

$$K = \sum_{i=1}^{n} h_i(X) - x \tag{12}$$

Let $\boldsymbol{\kappa} = [K_1, K_2, ..., K_n]$ be a vector whose generic element K_i represents the amount of capital distributed to the i^{th} participant, such that $\sum_{i=1}^{n} K_i = K$.

Following Clemente et al. (2024), we assume that the difference between the ex-ante contribution $h_i(X)$ of each participant i to the pool's overall risk and the ex-post capital allocation K_i received by the same participant remains constant across all participants. Formally, this condition is expressed as

$$h_i(X) - K_i = c,$$

where c is a constant. This approach defines an optimal allocation rule ensuring that each participant contributes equitably to the coverage of realized losses. Specifically, participants who contribute more to the pool's overall risk will bear a larger share of the realized losses or, conversely, will be entitled to a greater portion of the cashback distribution.

Clemente et al. (2024) formalized this idea through the following optimization problem:

$$\begin{cases} h_i(X) - K_i = c \\ \sum_{i=1}^{n} k_i = 1 \end{cases} \tag{13}$$

where the ex-post capital allocation K_i is expressed as a share of the total realized loss K, i.e., $K_i = k_i \cdot K$.

This formulation reflects the principle that the net position of each participant, measured as the difference between their initial contribution and the final capital allocation, should be identical across the group. In other words, all participants experience the same deviation, either gain or loss, from their expected contribution level, thereby ensuring fairness and symmetry in the redistribution process. This rule prevents any participant from being disproportionately rewarded or penalized relative to their ex-ante contribution and provides a transparent mechanism for sharing both surpluses and deficits arising from the pool's collective performance.

The solution to this problem yields the following optimal allocation:

$$\begin{cases} k_i^* = \dfrac{h_i(X) - \bar{x}}{\sum_{i=1}^n h_i(X) - x} \\ c = \bar{x} \end{cases} \tag{14}$$

where $\bar{x} = x/n$ represents the average realized cost per participant.

Considering Eq. (6), the optimal solution in Eq. (14) according to our contribution rule's framework becomes:

$$\begin{cases} k_i^* = \frac{\varphi_i(VaR_\alpha(Y)) + \frac{E(X)}{n} - \bar{x}}{VaR_\alpha(Y) + E(X) - x} \\ c = \bar{x} \end{cases} \tag{15}$$

Proof By substituting Eq. (12) into Eq. (13), we can rewrite it as:

$$\begin{cases} h_i(X) - k_i \cdot \left(\sum_{i=1}^n h_i(X) - x\right) = c \\ \sum_{i=1}^n k_i = 1 \end{cases} \tag{16}$$

Summing the first equation in system (16) over all $i = 1, \ldots, n$, we obtain:

$$\sum_{i=1}^{n} \left[h_i(X) - k_i \left(\sum_{j=1}^{n} h_j(X) - x \right) \right] = \sum_{i=1}^{n} c \tag{17}$$

$$\Rightarrow \sum_{i=1}^{n} h_i(X) - \left(\sum_{i=1}^{n} k_i \right) \left(\sum_{j=1}^{n} h_j(X) - x \right) = nc \tag{18}$$

Using the constraint $\sum_{i=1}^n k_i = 1$, we get:

$$\sum_{i=1}^{n} h_i(X) - \left(\sum_{j=1}^{n} h_j(X) - x \right) = nc \tag{19}$$

Hence:

$$c = \frac{x}{n} \tag{20}$$

□

3.4 Cost efficiency mechanisms in the P2P broker model

The proposed P2P insurance framework is broker-based and therefore relies on ex-ante contributions, in a manner similar to traditional insurance contracts. Policyholders pay a predefined premium, part of which is allocated to a mutual fund, while the remainder is transferred to a traditional insurer.

The potential cost efficiency of the P2P broker model does not arise mechanically from risk mutualization alone, but rather from a set of specific economic mechanisms embedded in its organizational structure. First, distribution and administrative costs may be reduced through fully digital intermediation by the broker platform, which replaces traditional agency networks and allows for lower marginal servicing costs. Second, the P2P arrangement may mitigate moral hazard and fraudulent behavior, as policyholders are directly exposed to the performance of the mutual fund through ex-post surplus redistribution mechanisms such as cashback payments or premium discounts. This shared exposure creates peer-monitoring incentives that can contribute to reductions in both claim frequency and claim severity.

Third, additional cost efficiencies may stem from improved risk selection. Participation in a P2P group is typically voluntary and often based on observable or self-declared characteristics, which can lead to the formation of relatively homogeneous pools and a reduction in adverse selection at the group level. Finally, the broker model relies on a predefined risk-sharing agreement with a traditional insurer, under which only losses exceeding a contractual threshold are transferred through an excess-of-loss structure. This allows the insurer to concentrate on tail-risk coverage, potentially improving capital efficiency and pricing accuracy.

It should be noted, however, that these cost-reducing effects are not guaranteed and depend critically on factors such as group size, pool homogeneity, governance rules, and the effectiveness of the platform's incentive mechanisms.

3.5 Frequency-Severity Model

We model individual losses using a standard *frequency-severity framework*, where each policyholder's total loss is represented as the sum of random claim amounts. Specifically, we define:

$$L_i = \sum_{h=1}^{NC_i} Z_{i,h}, \tag{21}$$

with $L_i = 0$ if $NC_i = 0$. Here:

- NC_i denotes the random number of claims for policyholder i;
- $Z_{i,h}$ denotes the cost (or severity) of the h-th claim for policyholder i.

We make the following assumptions:

- Conditional on policyholder i, the claim count NC_i is independent of the claim severities $Z_{i,h}$;
- Conditional on policyholder i, the severities $Z_{i,h}$ are independent and identically distributed (i.i.d.).

Under these assumptions, the moment generating function of L_i is given by:

$$M_{L_i}(t) = M_{NC_i}\left(\ln M_{Z_i}(t)\right), \tag{22}$$

where $M_{NC_i}(\cdot)$ and $M_{Z_i}(t)$ are the mgfs of the claim count and claim severity distributions for policyholder i, respectively.

From Eq. (22), the main cumulants of the distribution of L_i, and consequently of the P&L Y_i, can be derived analytically.

Modeling claim frequency. Following the literature (see, for example, Daykin et al. (1994)), we assume that the number of claims per individual follows a Poisson distribution:

$$NC_i \sim \text{Poisson}(f_i \cdot Q),$$

where:

- f_i is the expected claim frequency for policyholder i;
- Q is a structural (or contagion) variable capturing parameter uncertainty or common shocks affecting the claim count distribution. We assume that Q is defined only for positive values with $\mathbb{E}[Q] = 1$.

This framework introduces overdispersion through the mixing variable Q, which accounts for heterogeneity or contagion effects across policyholders.

Variance of individual losses. Under this specification, the standard deviation $\sigma(L_i)$ of the individual loss L_i is given by:

$$\sigma(L_i) = \sqrt{f_i \cdot a_{2,i} + f_i^2 \cdot a_{1,i}^2 \cdot \sigma_Q^2}, \tag{23}$$

where $a_{j,i}$ denotes the raw moment of order j of the severity distribution Z_i and σ_Q^2 is the variance of the contagion variable Q.

Covariance between policyholders. For two distinct policyholders i and j, with $i \neq j$, the covariance between their total losses is:

$$\begin{aligned} \text{Cov}(L_i, L_j) &= \mathbb{E}[\text{Cov}(L_i, L_j \mid Q)] + \text{Cov}(\mathbb{E}[L_i \mid Q], \mathbb{E}[L_j \mid Q]) \\ &= \text{Cov}(f_i \cdot a_{1,i} \cdot Q,\ f_j \cdot a_{1,j} \cdot Q) \\ &= \mathbb{E}[L_i] \cdot \mathbb{E}[L_j] \cdot \sigma_Q^2, \end{aligned} \tag{24}$$

This expression shows that the shared structural variable Q induces *positive dependence* among individual losses, leading to a natural form of correlation across policyholders within the risk pool.

From formulas (2), (23) and (24), we can express the variance of aggregate losses as:

$$\begin{aligned} \text{Var}(X) &= \sum_{i=1}^{n} \text{Var}(L_i) + 2 \sum_{1 \le i < j \le n} \text{Cov}(L_i, L_j) \\ &= \sum_{i=1}^{n} \left(f_i a_{2,i} + f_i^2 a_{1,i}^2 \sigma_Q^2 \right) + 2 \sum_{1 \le i < j \le n} \mathbb{E}[L_i]\mathbb{E}[L_j]\sigma_Q^2 \\ &= \sum_{i=1}^{n} f_i a_{2,i} + \sigma_Q^2 \left(\sum_{i=1}^{n} f_i a_{1,i} \right)^2 . \end{aligned}$$

Thus, the coefficient of variation of X is

$$\mathrm{cv}(X) = \frac{\sqrt{\sum_{i=1}^{n} f_i a_{2,i} + \sigma_Q^2 \left(\sum_{i=1}^{n} f_i a_{1,i} \right)^2}}{\sum_{i=1}^{n} f_i a_{1,i}} = \sqrt{\frac{\sum_{i=1}^{n} f_i a_{2,i}}{\left(\sum_{i=1}^{n} f_i a_{1,i}\right)^2} + \sigma_Q^2}. \tag{25}$$

This formulation (see last form in Eq. (25)) highlights two important components of risk. The first term represents the idiosyncratic risk from individual policyholders. As the portfolio size increases, this term decreases due to diversification, which is a standard result in risk pooling. The second term, represents non-diversifiable contagion risk, which cannot be mitigated by increasing the number of policyholders in the portfolio. This term remains constant regardless of the portfolio size and reflects the systemic risk shared by all policyholders.

To further analyze the behavior of the coefficient of variation (CV), we consider the case where the policyholders have identical severity distributions, i.e. $a_{j,i} = a_j \; \forall i$. Under this assumption, we can rewrite the CV as follows:

$$\mathrm{cv}(X) = \sqrt{\frac{a_2}{a_1^2 (\sum_{i=1}^{n} f_i)} + \sigma_Q^2} = \sqrt{\frac{1 + cv(Z)^2}{(\sum_{i=1}^{n} f_i)} + \sigma_Q^2}.$$

where $\mathrm{cv}(Z)$ is the coefficient of variation of the r.v. Z, assuming Z_i are identical for all policyholders.

It can be observed that when all policyholders share the same claim frequency, the portfolio benefits maximally from diversification. The coefficient of variation of aggregate claim amount decreases as the number of policyholders increases, reflecting the standard risk-pooling result that larger and more evenly distributed risks lead to lower relative volatility. When policyholders are heterogeneous in terms of claim frequency, the diversification benefit remains present but is weakened. In this case, the portfolio becomes more sensitive to large individual exposures, and the CV may decline more slowly as n increases. In extreme cases, a small number of large exposures may dominate the risk profile, so that the portfolio exhibits higher relative risk even as the number of policyholders grows. This highlights the importance of understanding the distribution of exposures within an insurance portfolio and the role of pool composition. Homogeneous portfolios allow for more effective risk pooling, whereas heterogeneous portfolios may require additional risk management strategies to address concentration risk arising from large exposures.

4 Numerical example

We apply the proposed methodology to a portfolio of motor other damage insurance policies. We assume the portfolio operates under a scheme that combines upfront contributions with a potential cashback mechanism. Individual policyholder losses are modeled using a frequency-severity framework, as defined in Eq. (21).

To model the total claim amount L_i for each policyholder, we introduce a systematic risk component represented by the common random variable Q, which affects all policyholders simultaneously.

Model Setup

- Number of policyholders: $n = 1,000$
- To generate the initial vector of individual claim frequencies, we draw one value for each policyholder from a Beta distribution. For each policyholder $i = 1, \dots, n$:

$$f_i \sim \text{Beta}(\alpha, \beta)$$

  with an average observed frequency between policyholders equal to 0.05 and a heterogeneity between policyholders based on a variance of 0.01.
  From these two conditions, we derive the Beta distribution parameters: $\alpha = 0.1875$ and $\beta = 3.5625$.
- The systematic risk component Q is modeled as:

$$Q \sim \text{Gamma}(\kappa, \theta)$$

  where κ is the shape parameter and θ is the scale parameter. In line with actuarial practice, we impose with $\mathbb{E}(Q) = \kappa \cdot \theta = 1$ and $\text{Var}(Q) = \kappa \cdot \theta^2 = 0.0025$. From these two conditions, we derive the Gamma distribution parameters: $\kappa = 400$ and $\theta = 0.0025$.
- The r.v. representing the number of claims for each policyholder, as previously discussed, is then obtained as:

$$NC_i \sim \text{Poisson}(f_i \cdot Q),$$

  where f_i denotes the expected frequency of claims for policyholder i, which is derived from a single simulation of the Beta distribution described earlier.
- The individual claim severity is modeled via i.i.d. LogNormal random variables. For each policyholder $i = 1, \dots, n$, and claim $h = 1, \dots, NC_i$:

$$Z_{i,h} \sim \text{LogNormal}\,(\mu_Z, \sigma_Z)$$

  with parameters μ_Z and σ_Z. Parameters are calibrated via method of moments, by assuming a mean equal to 2,000€ and standard deviation equal to 4,000€. The coefficient of variation $CV_Z = \frac{\sigma_Z}{\mu_Z}$ is therefore equal to 2, ensuring a reasonable level of skewness consistent with empirical insurance data.

Fair and pure premiums estimation

- The fair premium at the portfolio level is:

$$\mathbb{E}[X] = \mathbb{E}\left(\sum_{i=1}^{n} L_i\right) = \sum_{i=1}^{n} \mathbb{E}(L_i)$$

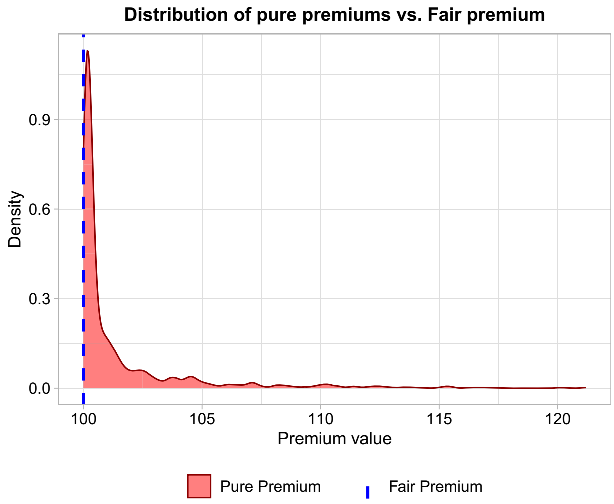

Fig. 1 Distribution of pure premiums $h_i(x)$ across policyholders versus fair premium $\frac{\mathbb{E}(X)}{n}$

By applying Eq. (21), we obtain:

$$\mathbb{E}(X) = \sum_{i=1}^{n} \mathbb{E}(NC_i) \cdot \mathbb{E}(Z_i) = \sum_{i=1}^{n} f_i \cdot \mathbb{E}(Z_i)$$

- Using the assumptions outlined above, the pure premium $h_i(X)$ for each policyholder can be computed through Eq. (8). The components involved in this calculation can be analytically derived by setting the desired confidence level α, and by applying the formulas in Eq. (23) and (24).
- To calculate K and k_i, it is necessary to obtain realizations of the aggregate claim amount x. This requires simulating the outcomes of the aggregate claims, which can be effectively done using Monte Carlo simulations of the compound process. These simulations are based on the framework developed earlier.

4.1 Results based on Gaussian assumption

Figure 1 displays the distribution of pure premiums $h_i(x)$, computed according to Eq. (8), based on Gaussian assumption and assuming a confidence level α of 5%, along with the fair premium $\frac{E(X)}{n}$. Relative to the average fair premium of €100, the implied safety loading coefficient ranges from approximately 0.13% to 21.2%, depending on the policyholder's risk profile. As discussed in Section 3, this safety loading is allocated ex ante based on each individual's marginal contribution to the portfolio's overall risk, as determined using the Shapley values within a VaR allocation framework. This results in a differentiated pricing scheme that accurately reflects each policyholder's exposure to tail risk.

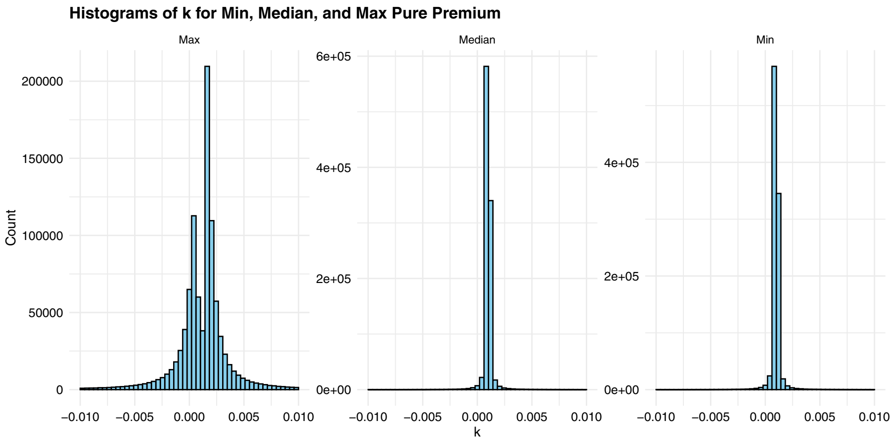

Fig. 2 Simulated distributions of the cashback percentage k_i^* for three representative policyholders, based on 1,000,000 different cashback scenarios. The three cases correspond to: the policyholder with the maximum value of $h_i(x)$ (left panel), the policyholder at the median of the $h_i(x)$ distribution (central panel), and the policyholder with the minimum value of $h_i(x)$ (right panel)

To quantify the expected cashback, we simulate the empirical distribution of aggregate payments at the portfolio level. As previously states, we assume that claim sizes follow a Lognormal distribution and simulate one million realizations of the total loss L_i for each policyholder using the compound mixed Poisson process defined in Eq. (21). This allows us to assess the variability of outcomes under the assumed risk structure.

Figure 2 shows the simulated distributions of the cashback percentage k_i^* for three representative policyholders, based on the simulated realizations of the aggregate claims amount. Specifically, we present results for the individuals with the minimum and maximum values of $h_i(x)$ - corresponding to the policyholders with the lowest and highest marginal risk contributions - as well as the policyholder at the median of the $h_i(x)$ distribution, representing a central case within the portfolio. The shape and spread of the k_i^* distributions differ across the selected individuals. For the policyholder with the highest pure premium $h_i(x)$, the distribution of k_i^* is noticeably wider, indicating greater variability in net contributions. This reflects their larger exposure to portfolio risk: in adverse scenarios, they are expected to contribute more, while in favorable scenarios they receive a higher refund. By contrast, the distribution is narrower for the low-risk policyholder. These results underscore the role of the Shapley-based allocation mechanism in linking premium contributions not only to expected claims but also to individual exposure to aggregate portfolio risk. In particular, the mechanism ensures that higher-risk policyholders bear a greater share of costs in unfavorable outcomes while remaining eligible for proportionally higher refunds when losses are lower than expected. This dynamic promotes fairness, financial sustainability, and stronger alignment between pricing and risk in mutualized insurance settings.

Figure 3 presents the distribution of cashback percentages k_i^* for all 1,000 policyholders under two distinct scenarios, based on 1,000,000 simulations. Specifically,

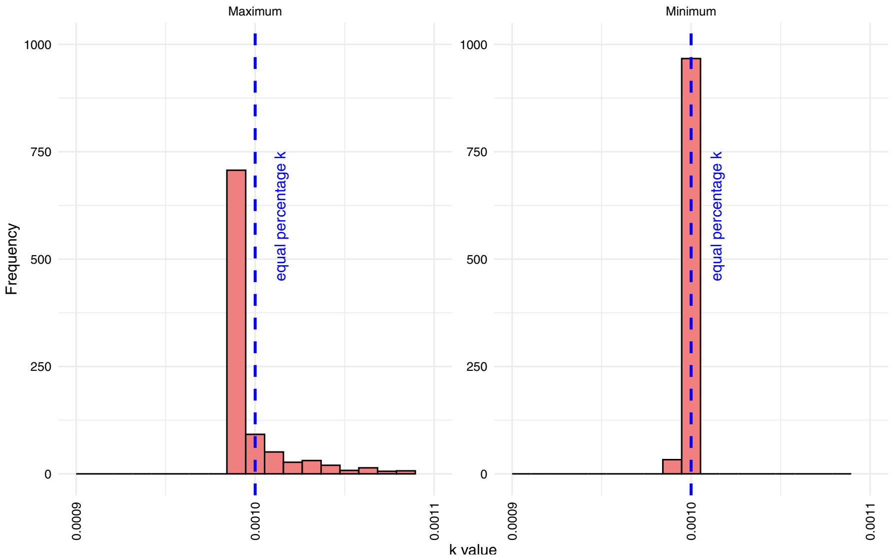

Fig. 3 Distributions of k_i^* for two policyholders under scenarios of maximum positive cashback (left panel) and minimum negative cashback (right panel)

we analyze the scenarios corresponding to the maximum and minimum values of the aggregate cashback amount K. In the scenario with the maximum aggregate cashback, the average cashback per policyholder is approximately 84€. Conversely, in the minimum cashback scenario, the average cashback is negative, around -1,588€, indicating a substantial loss. These two cases reveal markedly different behaviors. When the total aggregate amount x approaches zero (i.e., a very low payout), the distribution of the individual cashback percentages k_i^* closely resembles the distribution of the ex-ante contributions $h_i(x)$ (see the similarity between Figure 1 and left side of Figure 3). This suggests that, in favorable conditions, policyholders' cashback shares reflect their initial proportional contributions. On the other hand, when the total payout is very high, resulting in a negative cash-back, the influence of the ex-ante contributions on k_i^* diminishes significantly. In this scenario, all policyholders tend to share the loss equally, with each k_i^* approaching the uniform value $\frac{1}{n}$. This outcome highlights that, under extreme losses, the initial contribution proportions become less relevant, and the burden is distributed evenly across the portfolio.

In Figure 4, we examine the distribution of the variable $h_i(x)$ under various sensitivity scenarios involving the key model parameters. Starting from the base scenario, we consider the impact of doubling each parameter individually while keeping the others fixed. Specifically, we analyze the effects of doubling: (i) the average claim frequency $\mathbb{E}(f_i)$, (ii) the variance of the claim frequency, $\mathrm{Var}(f_i)$, (iii) the variance of the random variable Q, $\mathrm{Var}(Q)$, (iv) the expected average cost per claim, $\mathbb{E}(Z_i)$, (v) the coefficient of variation of the average cost, $CV(Z_i)$, and (vi) the number of policyholders, n. This comparative analysis highlights how each component influences the dispersion and shape of the $h_i(x)$ distribution. An increase in average parameters - such as the claim

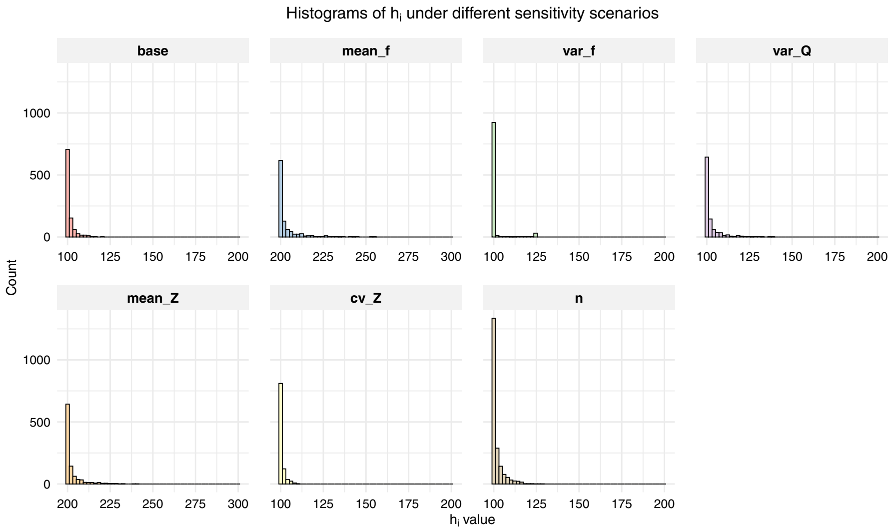

Fig. 4 Distributions of $h_i(x)$ under different sensitivity scenarios, illustrating the effects of doubling key model parameters, including the average claim frequency, variance of claim frequency, standard deviation of Q, expected average cost per claim, coefficient of variation of the average cost, and the number of policyholders

frequency or the average cost per claim - primarily results in a proportional increase in the average premium. Specifically, doubling these average values leads to an approximate doubling of the expected premium, without substantially affecting the shape or dispersion of the distribution of h_i. In contrast, the variability and overall shape of the h_i distribution are more sensitive to parameters that influence the heterogeneity in claim frequency across policyholders. Notably, when the variance of the claim frequency across policyholders is doubled, the variance of $h_i(x)$ increases by more than a factor of three. This indicates that the system responds non-linearly to changes in frequency heterogeneity. Moreover, when we increase the variance of the contagion parameter - capturing the variability in the common shock or cluster-level risk - the variance of $h_i(x)$ increases by almost four times. This significant rise is due to both greater individual-level volatility and enhanced dependence among policyholders induced by the contagion mechanism. Finally, increasing the number of policyholders results in a premium distribution that maintains the same average value, as the policyholders' characteristics are assumed to be similar across scenarios. However, the shape of the distribution undergoes some changes. While the overall dispersion of the $h_i(x)$ distribution tends to rise with the larger portfolio, reflecting the broader spread of risks, the relative volatility decreases due to diversification effects. The greater number of policyholders reduces the impact of individual variability, which in turn diminishes the skewness of the distribution. This phenomenon leads to a smoothing effect where idiosyncratic risks are more effectively averaged out across a larger pool, but it is partially mitigated by the presence of a systematic component Q that affects all the policyholders.

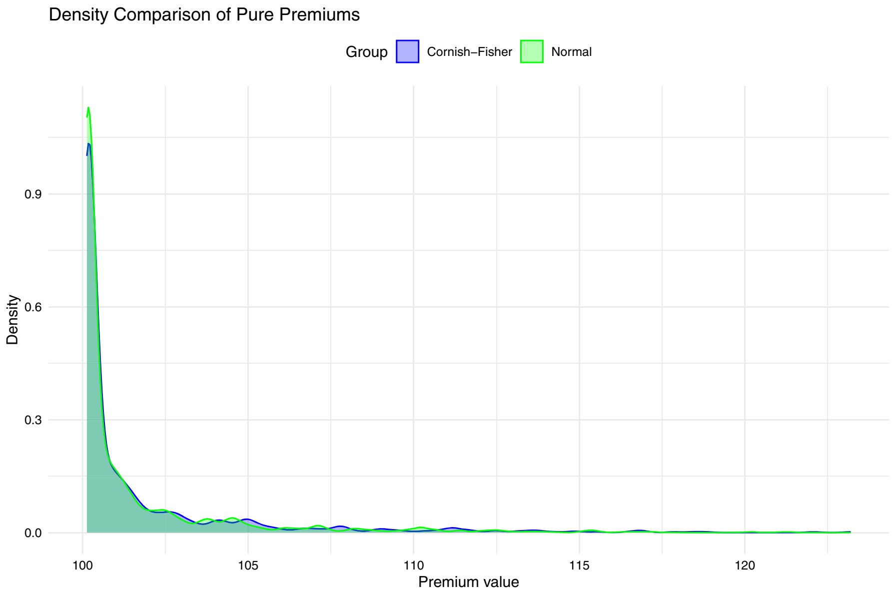

Fig. 5 Comparison of h_i where the Shapley value has been obtained using Gaussian assumption and Cornish-Fisher approximation

4.2 Results based on Cornish-Fisher approximation

In Figure 5, we present a comparison of pure premium distributions derived using a Shapley-based contribution, with one set of distributions assuming a Normal distribution (as specified in Eq. (8)) and the other applying the Cornish-Fisher approximation (outlined in Eq. (11)). The results reveal a close alignment between the two distributions, primarily due to the limited skewness of the portfolio's underlying distribution. This suggests that the Cornish-Fisher approximation, which is typically employed to account for higher-order moments (such as skewness and kurtosis) in cases with more extreme tails, has minimal impact in this instance.

The average pure premium, for example, is nearly identical across the two methods, with values of 101.46 and 101.60, respectively. Likewise, the maximum premium increases only slightly, from 121.19 under the Normal assumption to 123.20 when the Cornish-Fisher correction is applied. These similarities highlight that, while the Cornish-Fisher approximation provides a more refined adjustment to the distribution's tail, its effect is marginal when the underlying portfolio distribution exhibits low skewness. This suggests that for portfolios with relatively symmetric or slightly skewed distributions, the normal assumption may suffice for premium estimation without significant loss of accuracy. Consequently similar results have been obtained also in term of cashback, with the distribution slightly moved upwards due to the slight increase in the maximum pure premium.

In Figure 6, we present the distribution of the ex ante premium $h_i(x)$ under the same sensitivity scenarios previously analyzed, but with the Cornish-Fisher approximation applied to the Shapley-based contribution calculation (see Eq. (11)). Starting with the base scenario, we assess the impact of doubling each parameter individually while holding the others constant. As before, we analyze the effects of doubling: (i) the average claim frequency $\mathbb{E}(f_i)$, (ii) the variance of the claim frequency, $\text{Var}(f_i)$, (iii) the variance of the random variable Q, $\text{Var}(Q)$, (iv) the expected average cost per claim, $\mathbb{E}(Z_i)$, (v) the coefficient of variation of the average cost, $CV(Z_i)$, and (vi) the number of policyholders, n. This comparative analysis highlights how each parameter influences the dispersion and shape of the $h_i(x)$ distribution.

The main behaviors observed are consistent with the previous case of the Normal distribution, but some important differences emerge when applying the Cornish-Fisher approximation. Again, an increase in average parameters, such as the claim frequency or the average cost per claim, results in a proportional increase in the average premium, without substantially affecting the shape or dispersion of the distribution of h_i. This observation reinforces the conclusion that the normal assumption can be sufficient for premium estimation without significant loss of accuracy. Moreover, increasing the number of policyholders reduces the impact of individual variability, smoothing the overall distribution. Additionally, as the skewness of the distribution is reduced, the Cornish-Fisher approximation and the Normal assumption become closer to one another in terms of their resulting shapes.

In contrast, the variance of the random variable Q produces a more pronounced increase in the volatility of $h_i(x)$ compared to the Gaussian case. This can be attributed to the fact that a higher variance also leads to increased skewness in the distribution of Q, which, in turn, introduces more negative skewness into the profit and loss distribution Y. This highlights the differences between the Cornish-Fisher approximation and the Normal assumption, especially when the variance of Q is large.

Another key difference becomes more apparent in the effect of the coefficient of variation ($CV(Z_i)$) of the claim size distribution. Doubling $CV(Z_i)$ leads to a significantly more skewed profit and loss distribution. For distributions of the claim size with very high volatility or fat tails, the Gaussian assumption tends to fail by underestimating the Shapley-based contribution. Specifically, when the Cornish-Fisher approximation is used, the maximum value of $h_i(x)$ increases, reaching 147, compared to 111 when the Normal assumption is applied. This underscores the limitations of the Gaussian approach in capturing the full impact of extreme volatility or skewness in the distribution of claim sizes.

5 Concluding remarks

This paper proposes an ex-ante contribution mechanism for P2P insurance schemes, based on the Shapley value applied to the VaR of the difference between the expected and realized total loss of the group. By combining cooperative game theory with probabilistic modeling, the mechanism ensures that contributions are allocated according to each participant's marginal contribution to risk. This approach maintains simplicity, as the total expected loss is equally shared and deviations from expectations are proportionally distributed.

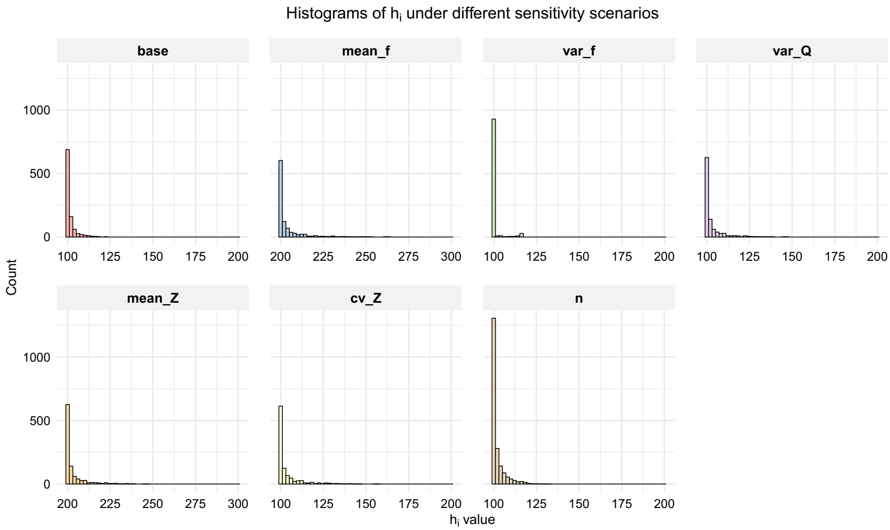

Fig. 6 Distributions of $h_i(x)$ under different sensitivity scenarios, but with the Cornish-Fisher approximation applied to the Shapley-based contribution calculation (see formula (11)). Figure shows the effects of doubling key model parameters, including the average claim frequency, variance of claim frequency, variance of Q, expected average cost per claim, coefficient of variation of the average cost, and the number of policyholders

The proposed framework advances the literature on capital allocation and contribution mechanisms in P2P insurance by linking risk-sensitive fairness with practical applicability. It offers several advantages. It enhances transparency in contribution rules, aligns incentives among participants, and can be integrated into existing broker-based P2P schemes without requiring fundamental redesign.

Nevertheless, some limitations remain. First, reliance on the VaR measure entails sensitivity to tail events and may underestimate systemic risks under extreme scenarios. Second, the model does not explicitly account for behavioral aspects such as trust, cooperation, and reputation, that play a central role in the success of P2P arrangements.

Future research could extend the proposed mechanism along several directions. On the methodological side, replacing VaR with coherent risk measures such as Conditional Value-at-Risk (CVaR) could improve robustness and sensitivity to extreme outcome. On the institutional side, embedding the mechanism within regulatory sandbox environments would provide valuable insights into its operational feasibility. Finally, integrating reputation systems or blockchain-based transparency tools could help bridge the gap between formal allocation rules and the behavioral dynamics governing peer communities.

To conclude, the Shapley-based ex-ante contribution mechanism provides a rigorous yet practical tool for structuring P2P insurance, striking a balance between simplicity and fairness. Its ultimate success will depend on its adaptability to real-world platforms and its compatibility with evolving regulatory frameworks. At present, regulation in this area remains fluid; clear and proportionate rules are needed to address solvency, governance, and consumer protection concerns, thereby supporting the insurance industry's ability to build sustainable and effective P2P models. Within this context, the contribution and allocation mechanism proposed in this paper should be interpreted as a tool for improving transparency in risk sharing, rather than as a substitute for regulatory protection. Although Shapley-based rules can reduce uncertainty by aligning contributions with marginal risk and limiting cross-subsidization, regulatory oversight remains essential to address solvency, governance, and consumer protection concerns in P2P insurance markets.

Acknowledgements Gian Paolo Clemente acknowledges funding by the European Union – Next Generation EU. Project PRIN 2022 "Building resilience to emerging risks in financial and insurance markets". Project code: 2022FWZ2CR - CUP J53D23004560008. The views and opinions expressed are only those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Author Contributions All authors contributed significantly to all stages of the work, from conceptualization to the final version. All authors have read and agreed to the published version of the manuscript.

Funding Open access funding provided by Università Cattolica del Sacro Cuore within the CRUI-CARE Agreement.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Non-financial interests The authors have no relevant financial or non-financial interests to disclose.

References

Abdikerimova, S., Feng, R.: Peer-to-peer multi-risk insurance and mutual aid. Eur. J. Oper. Res. **299**, 735–749 (2022). https://doi.org/10.1016/j.ejor.2021.09.017

Acerbi, C.: Spectral measures of risk: A coherent representation of subjective risk aversion. Journal of Banking & Finance **26**(7), 1505–1518 (2002)

Bellini, F., Klar, B., Müller, A., Rosazza Gianin, E.: Generalized quantiles as risk measures. Insurance Math. Econom. **54**, 41–48 (2014)

Charpentier, A., Kouakou, L., Lowe, M., Ratz, P., Vermet, F.: Collaborative insurance sustainability and network structure. (2021) arXiv preprint arXiv:2107.02764

Chen, X., Hu, Z., Wang, S.: Stable Risk Sharing and Its Monotonicity (June 15, 2017). (2017) Available at SSRN: https://ssrn.com/abstract=2987631 or http://dx.doi.org/10.2139/ssrn.2987631

Clemente, G.P., Marano, P.: The broker model for peer-to-peer insurance: An analysis of its value. The Geneva Papers on Risk and Insurance - Issues and Practice **45**(2), (2020)

Clemente, G.P., Levantesi, S., Piscopo, G.: Optimal cashback in a cooperative framework for peer-to-peer insurance coverages. Ann. Oper. Res. (2023). https://doi.org/10.1007/s10479-023-05687-7

Clemente, G.P., Levantesi, S., Piscopo, G.: Risk Sharing Rule and Safety Loading in a Peer to Peer Cooperative Insurance Model. Decisions Econ. Finan. (2024). https://doi.org/10.1007/s10203-024-00438-0

Colini-Baldeschi, R., Scarsini, M., Vaccari, S.: Variance Allocation and Shapley Value. Methodology Computing in Applied Probability **20**, 919–933 (2018). https://doi.org/10.1007/s11009-016-9540-5

Cornish, E.A., Fisher, R.A.: Moments and cumulants in the specification of distributions. Revue de l'Institut international de Statistique 307–320 (1938)

Daykin, C., Pentikáinen, T., Pesonen, M.: Practical Risk Theory for Actuaries. Chapman & Hall (1994)

Denuit, M., Dhaene, J.: Convex order and comonotonic conditional mean risk sharing. Insurance Math. Econom. **51**, 265–270 (2012)

Denuit, M.: Size-biased transform and conditional mean risk sharing, with application to P2P insurance and tontines. ASTIN Bulletin **49**, 591–617 (2019)

Denuit, M.: Size-biased risk measures of compound sums. North American Actuarial Journal (2020)

Denuit, M., Robert, C.Y.: Large Loss Behaviour of Conditional Mean Sharing. ASTIN Bulletin **50**(3), 1093–1122 (2020)

Denuit, M., Robert, C.Y.: Collaborative Insurance with Stop-Loss Protection and Team Partitioning. North American Actuarial Journal, 1-18. (2021) Available at: http://dx.doi.org/10.1080/10920277.2020.1855199

Denuit, M., Robert, C.Y.: From risk sharing to pure premium for a large number of heterogeneous losses. Insurance Math. Econom. **96**, 116–126 (2021)

Denuit, M., Dhaene, J., Robert, C.Y.: Risk-sharing rules and their properties, with applications to peer-to-peer insurance. Journal of Risk and Insurance **89**(3), 615–667 (2022)

Dhaene, J., Tsanakas, A., Vanduffel, S.: On Distortion Risk Contribution Measures. Insurance Math. Econom. **51**(3), 533–541 (2012). https://doi.org/10.1016/j.insmatheco.2012.03.010

European Insurance and Occupational Pensions Authority (EIOPA) (2019). Peer-to-peer insurance: Supervisory statement on licensing requirements. Frankfurt: EIOPA. https://www.eiopa.europa.eu

European Insurance and Occupational Pensions Authority (EIOPA) (2020). Digitalisation in insurance and pensions. Frankfurt: EIOPA. https://www.eiopa.europa.eu

European Insurance and Occupational Pensions Authority (EIOPA) (2023). Supervisory convergence in digitalisation. Frankfurt: EIOPA. https://www.eiopa.europa.eu

Feng, R., Liu, M., Zhang, N.: A unified theory of decentralized insurance. Insurance Math. Econom. **119**, 157–178 (2024). https://doi.org/10.1016/j.insmatheco.2024.08.008

Feng, R., Liu, M., Taylor, S.: Peer-to-peer risk sharing with an application to flood risk pooling. Ann. Oper. Res. **321**, 813–842 (2023). https://doi.org/10.1007/s10479-022-04841-x

Föllmer, H., Schied, A.: Stochastic finance: An introduction in discrete time (4th ed.). De Gruyter (2016)

Galeotti, M., Rabitti, G.: Tail variance allocation, Shapley value, and the majorization problem. J. Appl. Probab. **61**(1), 153–171 (2024). https://doi.org/10.1017/jpr.2023.28

Global InsurTech Report (2023). Peer-to-peer insurance: Market outlook and challenges. London: Deloitte Insights. https://www2.deloitte.com

Klugman, S.A., Panjer, H.H., Willmot, G.E.: Loss Models: From Data to Decisions, 5th edn. Wiley, Hoboken (2019)

Hardy, M.R.: Approximating the aggregate claims distribution. Encyclopedia of Actuarial Science, John Wiley & Sons, Ltd (2006)

Hagan, P.S., Lesniewski, A., Skoufis, G.E., et al.: Portfolio risk allocation through Shapley value. International Journal of Financial Engineering **2350004**,(2023). https://doi.org/10.1142/S2424786323500044

Kaas, R., Goovaerts, M., Dhaene, J., Denuit, M.: Modern Actuarial Risk Theory: Using R. Springer (2008)

Landsberger, M., Meilijson, I.: A Class of Cooperative Games Continuous in the Individual Risk. Insurance Math. Econom. **25**(1), 97–106 (1999). https://doi.org/10.1016/S0167-6687(99)00010-2

Levantesi, S., Piscopo, G.: Mutual peer-to-peer insurance: The allocation of risk. Journal of Co-operative Organization and Management **10**(1), 100154 (2022). https://doi.org/10.1016/j.jcom.2021.100154

Mango, D.F.: An Application of Game Theory: Property Catastrophe Risk Load. Proceedings of the Casualty Actuarial Society **85**, 157–86 (1998)

National Association of Insurance Commissioners (NAIC) (2021). Peer-to-peer insurance: Innovation and regulatory considerations. Washington, DC: NAIC. https://content.naic.org

Ostrowska, A., Ziemiak, M.: Peer-to-peer insurance: Legal and regulatory challenges. Journal of European Consumer and Market Law **9**(4), 163–172 (2020). (https://kluwerlawonline.com)

Rego, L.-M., Carvalho, C.J.: Insurance in today's sharing economy: new challenges ahead or a return to the origins of insurance? In: Marano, P., Noussia, K. (eds.) InsurTech: A Legal and Regulatory View, pp. 27–47. Springer, Berlin (2019)

Roth, A.E.: The Shapley value. Essays in honor of L. S. Shapley (1988), Cambridge University Press, Cambridge (1988)

Shapley, L.S.: A Value for n-Person Games. In Contributions to the Theory of Games, Princeton: Princeton University Press. Vol. II (1953)

Sun, R.-T., Garimella, A., Han, W., Chang, H.-L., Shaw, M.J.: Transformation of the Transaction Cost and the Agency Cost in an Organization and the Applicability of Blockchain - A Case Study of Peer-to-Peer Insurance. Frontiers in Blockchain **3**, 24 (2020). https://doi.org/10.3389/fbloc.2020.00024

Wang, S.S.: A class of distortion operators for pricing financial and insurance risks. J. Risk Insur. **67**(1), 15–36 (2000)