CAUCHY – Jurnal Matematika Murni dan Aplikasi
Volume 11 (2) (2026), Pages 2121–2135
p-ISSN: 2086-0382; e-ISSN: 2477-3344

Copula Regression for Deductible–Claim Dependence in Non-Life Insurance: A Stratified Evaluation

Dwi Mifta Mahanani[1], Feby Indriana Yusuf*[1], Dzaki Ferlian Nugroho[1], and Tuti Sariningsih Budi Utami[2]

[1]*Department of Mathematics, Universitas Brawijaya, Malang, Indonesia*
[2]*Ministry of Finance, Jakarta, Indonesia*

Abstract

Deductible choice may be statistically associated with claim outcomes after observed rating variables have been taken into account. The contractual deductible is distinguished from the normalized deductible ratio used as the continuous working margin, while claim count remains discrete and individual severity remains continuous. A Canonical Vine pair-copula construction links the deductible ratio, claim frequency, and individual claim severity. The application uses 1,038 entities from the Wisconsin Local Government Property Insurance Fund observed from 2006 to 2010; models are estimated on 2006-2009 data and evaluated on the 2010 temporal holdout. The estimated dependence is negative for deductible-frequency, positive for deductible observed severity, and negative for frequency-severity conditional on the deductible. These patterns are interpreted as conditional associations that may be consistent with selection and claim-reporting mechanisms, rather than as causal evidence of adverse selection or moral hazard. The dependence model attains a higher observation-level logarithmic score than the exogenous-deductible benchmark for 64.55% of frequency observations and 71.09% of severity observations. Its global Insurance Gini improvement is 22.54 points, with positive point estimates across all five deductible strata but greater uncertainty in the highest-deductible group. The results show that deductible information can improve out of sample risk discrimination, while also demonstrating the importance of support assumptions, benchmark definitions, and cautious behavioral interpretation.

Keywords: Canonical Vine; compound distribution; copula regression; deductible choice; Insurance Gini; non-life insurance.

1. Introduction

Accurate non-life insurance ratemaking requires an understanding of how observable policyholder characteristics, contract choices, and claim outcomes are related. A deductible allocates the first layer of a loss to the insured and may also convey information about risk retention and reporting behavior. Consequently, the deductible may remain statistically associated with claim frequency or observed claim severity after conventional rating variables have been controlled for. Treating that association as absent may weaken pricing and risk segmentation. Nevertheless, residual association is not equivalent to causal identification: adverse selection, reporting thresholds, omitted risk characteristics, contract design, and other mechanisms can generate similar empirical patterns [1, 2].

***Corresponding author**. E-mail: febyindrianay@ub.ac.id

Received: June 10, 2026 Revised: August 29, 2026 Accepted: September 17, 2026
DOI: https://doi.org/10.18860/cauchy.v11i2.43524

Frequency and severity are also not necessarily independent. Dependence can arise from unobserved heterogeneity, repeated observations of the same policyholder, and differences in claim-reporting practices. Frequency–severity regression and longitudinal insurance studies therefore emphasize the value of joint rather than separate modeling [3–5]. Copulas are useful in this setting because they separate marginal regression specifications from dependence assumptions. Pair-copula constructions further decompose a multivariate distribution into interpretable bivariate building blocks and can accommodate combinations of continuous and discrete outcomes [6–8].

The methodological foundation of the present application is the copula regression model for compound distributions with endogenous covariates introduced by Shi and Lee [9]. That study jointly modeled a policyholder's normalized deductible, number of claims, and individual claim amounts and applied the model to the Wisconsin Local Government Property Insurance Fund (LGPIF). The present article does not claim the original construction of that architecture. Instead, it reproduces the core specification and extends its evaluation by examining whether the predictive and discriminatory gain associated with deductible–claim dependence is stable across contractually meaningful deductible strata. This transparent positioning is important because the same LGPIF portfolio and the publicly documented computational framework are used.

Three issues receive particular attention. First, the contractual deductible D is a finite-menu choice, whereas the primary regression uses the normalized ratio $R = D/C$, where C is coverage. Variation in coverage gives R a substantially richer empirical support, allowing it to be treated as a positive continuous working margin; claim count remains discrete. A discrete/ordinal treatment of D is used as a specification check, for which the selected dependence directions and substantive conclusions remain unchanged. Second, because the same 1,038 entities are followed from 2006 to 2010, the analysis is described as a pooled, stage-wise estimation with a temporal holdout rather than as an independent cross-section. Third, dependence parameters are interpreted as conditional associations, not as causal estimates of adverse selection or moral hazard. Recent panel extensions provide an important direction for a fully dynamic treatment of deductible choice and claims [10].

The empirical results show negative deductible–frequency dependence, positive deductible–observed-severity dependence, and negative frequency–severity dependence conditional on the deductible. The dependence model produces higher observation-level logarithmic scores than an exogenous-deductible benchmark for 64.55% of frequency observations and 71.09% of severity observations. The global Insurance Gini comparison favors the dependence model by 22.54 Gini points. Positive point estimates are also observed in each deductible stratum, although the highest-deductible group has a wide confidence interval and is not statistically distinguishable from zero at the 5% level.

The contribution of this study is therefore an out-of-sample and deductible-stratified assessment of an established copula pricing framework. It clarifies the mixed-support likelihood, distinguishes exogenous-deductible and total-independence benchmarks, documents the stage-wise estimation and temporal validation design, and evaluates heterogeneity in risk discrimination across deductible groups. The remainder of the paper is organized as follows. Section 2 presents the model and estimation procedure. Section 3 reports the empirical findings and out-of-sample validation. Section 4 discusses risk discrimination and Insurance Gini results. Section 5 concludes.

2. Methodology and Model Specification

This study adopts the compound-distribution copula regression architecture of Shi and Lee [9] and evaluates its performance across deductible strata. The term *deductible–claim dependence* is used for residual statistical dependence after conditioning on observed rating variables. It does not, by itself, identify a causal effect of the deductible or establish a particular asymmetric-information mechanism.

2.1. Observation unit, support, and working assumptions

Let $i = 1, \ldots, m$ index insured entities and let $t = 1, \ldots, T_i$ index policy years. The observed record for entity i in year t is

$$\mathcal{O}_{it} = (D_{it}, C_{it}, R_{it}, N_{it}, Y_{it1}, \ldots, Y_{itN_{it}}, \mathbf{x}_{it}), \tag{1}$$

where D_{it} is the contractual deductible, C_{it} is coverage, $R_{it} = D_{it}/C_{it}$ is the normalized deductible ratio, $N_{it} \in \{0, 1, 2, \ldots\}$ is claim count, $Y_{itj} > 0$ is the ground-up amount of the jth reported claim, and $\mathbf{x}_{it}$ contains the observed rating variables.

The contractual deductible D_{it} is selected from a finite menu. The primary model nevertheless treats R_{it} as a positive continuous working margin because variation in coverage produces a much richer support for D_{it}/C_{it} than for the nominal deductible itself. Claim count remains discrete and individual severity remains continuous. This distinction resolves the apparent inconsistency between a finite deductible menu and differentiation with respect to r. Copulas involving discrete variables are not unique outside their attainable support; consequently, interpretation is restricted to the fitted parametric representation and its predictive implications [8, 11]. As a support sensitivity check, the analysis was repeated using the ordered contractual deductible rather than the continuous ratio. The signs of the dependence estimates, selected families, and substantive predictive conclusions were unchanged.

Conditional on $(R_{it}, N_{it}, \mathbf{x}_{it})$, the individual severities within a policyholder-year are assumed identically distributed and conditionally independent:

$$Y_{it1}, \ldots, Y_{itN_{it}} \overset{\text{c.i.i.d.}}{\sim} F_Y(\,\cdot \mid R_{it}, N_{it}, \mathbf{x}_{it}). \tag{2}$$

This working assumption makes the random dimension of the severity vector explicit. When $N_{it} = 0$, no severity is observed and the associated product in the likelihood is defined as one. Residual dependence among claims from the same entity or year is not separately parameterized.

For a generic reported claim, define

$$F_{RNY}(r, n, y \mid \mathbf{x}) = \Pr(R \le r, N \le n, Y \le y \mid \mathbf{x}). \tag{3}$$

Let $u_R = F_R(r \mid \mathbf{x})$, $u_Y = F_Y(y \mid \mathbf{x})$, $u_N^+ = F_N(n \mid \mathbf{x})$, and $u_N^- = F_N(n-1 \mid \mathbf{x})$, with $F_N(-1 \mid \mathbf{x}) = 0$. For continuous R and Y and discrete N, the joint mass–density contribution is

$$f_{RNY}(r, n, y \mid \mathbf{x}) = \frac{\partial^2}{\partial r\, \partial y} \left[F_{RNY}(r, n, y \mid \mathbf{x}) - F_{RNY}(r, n-1, y \mid \mathbf{x}) \right]. \tag{4}$$

Thus, a finite difference is applied to the count coordinate, whereas derivatives are applied only to the continuous coordinates.

2.2. C-vine pair-copula representation

The normalized deductible is used as the root of a three-dimensional C-vine. Let C_{RN} link the deductible ratio and count, C_{RY} link the deductible ratio and severity, and $C_{NY|R}$ link the count and severity after conditioning on the deductible ratio. This is the standard pair-copula decomposition described by Aas et al. [6] and adapted to insurance mixed margins by Shi and Yang [7] and Shi and Lee [9].

For the continuous–discrete pair (R, N), define

$$H_N^{\pm}(r, n \mid \mathbf{x}) = \frac{\partial}{\partial u_R} C_{RN}\Big(u_R, u_N^{\pm}\Big). \tag{5}$$

The conditional count probability is then

$$p_{N|R}(n \mid r; \mathbf{x}) = H_N^+(r, n \mid \mathbf{x}) - H_N^-(r, n \mid \mathbf{x}). \tag{6}$$

This finite-difference expression is required because N is discrete.

For the continuous pair (R, Y),

$$f_{Y|R}(y \mid r; \mathbf{x}) = f_Y(y \mid \mathbf{x}) c_{RY}(u_R, u_Y), \qquad c_{RY}(u_R, u_Y) = \frac{\partial^2 C_{RY}(u_R, u_Y)}{\partial u_R \, \partial u_Y}, \tag{7}$$

and its conditional distribution is

$$V_Y = F_{Y|R}(y \mid r; \mathbf{x}) = \frac{\partial}{\partial u_R} C_{RY}(u_R, u_Y). \tag{8}$$

Set $V_N^+ = H_N^+(r, n \mid \mathbf{x})$ and $V_N^- = H_N^-(r, n \mid \mathbf{x})$. Let $\partial_2 C_{NY|R}$ denote differentiation with respect to the second argument. The conditional severity density given both $R = r$ and $N = n$ is

$$\begin{aligned} f_{Y|RN}(y \mid r, n; \mathbf{x}) = f_{Y|R}&(y \mid r; \mathbf{x}) \\ &\times \frac{\partial_2 C_{NY|R}(V_N^+, V_Y) - \partial_2 C_{NY|R}(V_N^-, V_Y)}{V_N^+ - V_N^-}. \end{aligned} \tag{9}$$

Equation (9) combines differentiation in the continuous severity coordinate with a finite difference across the attainable count interval. The implementation adopts the simplifying-vine assumption: conditional copula parameters do not vary with the realized value of R after the conditional probability transforms have been formed.

Under the conditional-independence assumption in Eq. (2), the policyholder-year likelihood contribution becomes

$$L_{it} = f_R(r_{it} \mid \mathbf{x}_{it}) p_{N|R}(n_{it} \mid r_{it}; \mathbf{x}_{it}) \prod_{j=1}^{n_{it}} f_{Y|RN}(y_{itj} \mid r_{it}, n_{it}; \mathbf{x}_{it}). \tag{10}$$

This formulation accommodates the absence of severity observations when the count is zero and avoids treating the severity vector as a fixed-dimensional random variable.

2.3. Dependence and independence benchmarks

Three specifications are distinguished. The *full dependence model* estimates C_{RN}, C_{RY}, and $C_{NY|R}$. The primary benchmark is the *exogenous-deductible model*, in which C_{RN} and C_{RY} are set to independence copulas while dependence between frequency and severity is retained through an unconditional C_{NY}. Its policyholder-year contribution is

$$L_{it}^{\text{exo}D} = f_R(r_{it} \mid \mathbf{x}_{it}) p_N(n_{it} \mid \mathbf{x}_{it}) \prod_{j=1}^{n_{it}} f_{Y|N}(y_{itj} \mid n_{it}; \mathbf{x}_{it}). \tag{11}$$

This benchmark removes only the incremental predictive information in the deductible. A separate *total-independence model* sets all three pair copulas to independence. The terms "exogenous-deductible" and "total independence" are therefore not used interchangeably.

2.4. Stage-wise estimation and panel interpretation

The parameters are estimated using a staged Inference Functions for Margins procedure rather than simultaneous full-information maximum likelihood. Let $\boldsymbol{\theta}_1$ collect marginal parameters, $\boldsymbol{\theta}_2$ collect the deductible–frequency copula parameters, and $\boldsymbol{\theta}_3$ collect the severity and remaining copula parameters. The pooled working log-likelihood is

$$\ell(\boldsymbol{\theta}) = \sum_{i=1}^{m} \sum_{t=1}^{T_i} \log L_{it}(\boldsymbol{\theta}). \tag{12}$$

The parameter blocks are obtained sequentially:

$$\begin{aligned}
\widehat{\boldsymbol{\theta}}_1 &= \arg\max_{\boldsymbol{\theta}_1} \ell_1(\boldsymbol{\theta}_1), \\
\widehat{\boldsymbol{\theta}}_2 &= \arg\max_{\boldsymbol{\theta}_2} \ell_2(\widehat{\boldsymbol{\theta}}_1, \boldsymbol{\theta}_2), \\
\widehat{\boldsymbol{\theta}}_3 &= \arg\max_{\boldsymbol{\theta}_3} \ell_3(\widehat{\boldsymbol{\theta}}_1, \widehat{\boldsymbol{\theta}}_2, \boldsymbol{\theta}_3).
\end{aligned} \tag{13}$$

The first stage estimates the marginal regressions. The second estimates the deductible–frequency dependence from the fitted probability transforms. The final stage estimates the severity margin, deductible–severity copula, and conditional frequency–severity copula while retaining the earlier estimates. Standard errors reported for the copula table are conditional observed-information standard errors from their corresponding stage; they do not represent a full simultaneous-MLE covariance matrix. This distinction is relevant because first-stage uncertainty and serial dependence are not completely propagated by the final-stage Hessian [12].

Although Eq. (12) uses policyholder-year contributions, observations from the same entity are not conceptually independent. The point estimates should therefore be interpreted as pooled working-likelihood estimates for a short panel. The analysis does not claim to fit an entity-specific random effect or a dynamic panel copula. This limitation motivates comparison with recent panel methods for insurance contracts with endogenous deductibles [10].

Numerical optimization uses the bounded `nlminb` optimizer in R for all custom likelihoods. The starting values are deterministic and are obtained from simpler fitted models: the deductible-ratio GB2 regression starts from the coefficients of an OLS regression for $\log(R)$ with $(\sigma, \alpha_1, \alpha_2) = (1, 1, 1)$; the zero-one-inflated negative-binomial frequency model starts from the corresponding zero-inflated negative-binomial fit, with the one-inflation coefficients initialized at zero; the severity GB2 regression starts from an OLS regression for $\log(Y)$ with $(\sigma, \alpha_1, \alpha_2) = (1, 1, 1)$; the (R, N) Student-t copula starts from $(\rho, \nu) = (0, 5)$; and the final severity/copula stage starts from the fitted severity-margin estimates together with $(0.5, -1.5, 5, 5)$ for the two copula parameters and their degrees-of-freedom arguments. The final joint optimization imposes the parameter bounds used by the copula family, including $\rho \in [-50, 50]$ for the first copula and $\theta_{\text{signed}} \in [-50, -1.001]$ for the rotated Gumbel parameterization; the (R, N) Student-t copula uses $\rho \in [-0.999, 0.999]$ and $\nu \geq 1$.

The baseline custom optimizations use at most 500 objective-function evaluations; the deductible-ratio fit also uses a maximum of 500 iterations. Unless explicitly overridden, `nlminb`'s default convergence tolerances are retained: relative function tolerance 10^{-10}, parameter tolerance 1.5×10^{-8}, and false-convergence tolerance 2.2×10^{-14}. Re-estimation with evaluation limits of 1,000 and 2,000 returns the same objective values and parameter estimates to the reported precision. These checks support numerical stability for the implemented starting configuration, but they are not presented as proof of a global optimum. Numerical Hessians are computed with `numDeriv::hessian` at the final estimates, and the reported conditional standard errors use the inverse observed Hessian from the corresponding estimation stage.

2.5. Model selection and temporal validation

Copula families are compared using the negative log-likelihood, Akaike information criterion, and Bayesian information criterion,

$$\text{AIC} = 2\,\text{NLL} + 2k, \qquad \text{BIC} = 2\,\text{NLL} + k \log n, \tag{14}$$

where k is the number of estimated parameters and n is the estimation sample size relevant to the stage. Lower values indicate a preferred candidate within the evaluated set.

The model is estimated using policy years 2006–2009 and evaluated on the 2010 holdout. Because the holdout may contain entities observed during training, this is a temporal one-year-ahead assessment rather than validation on previously unseen entities. Predictive distributions are evaluated with the logarithmic score [13, 14]:

$$S(p, w) = \log p(w), \tag{15}$$

where a larger value is better. In addition to the total log score, the empirical win rate of the full model relative to the exogenous-deductible benchmark is

$$\widehat{\pi}_{\text{win}} = \frac{1}{n_{\text{test}}} \sum_{q=1}^{n_{\text{test}}} I\left\{ S_q^{\text{full}} \geq S_q^{\text{exo}D} \right\}. \tag{16}$$

This quantity is a sample proportion of holdout observations, not a classification accuracy or a model-based probability.

3. Results and Discussion

This section presents and discusses the main empirical findings. We begin by describing the data and identifying their key exploratory patterns. Next, we examine the fitted marginal distributions and the dependence structure captured by the copula models. Finally, we evaluate the out-of-sample predictive performance and discuss the implications of the results.

3.1. Data and exploratory patterns

The empirical application uses the Wisconsin Local Government Property Insurance Fund data analyzed by Shi and Lee [9]. The Fund provides building and contents coverage to local government entities in Wisconsin. The balanced panel contains 1,038 entities observed annually from 2006 to 2010, yielding 5,190 policyholder-year observations. The first four years are used for estimation and 2010 is retained for temporal validation. Available variables include the contractual deductible, coverage, claim count, per-occurrence ground-up loss, and entity classification.

Table 1: Description and descriptive statistics of outcome and explanatory variables

Variable	Description	Mean	SD
Deductible	Amount of deductible per million dollars of coverage	1,982.787	41,194.260
ClaimFreq	Number of claims per policyholder-year	0.589	1.537
ClaimSev	Amount of an individual positive claim in dollars	19,515.000	154,221.200
Entity type indicators			
City	1 for a city entity and 0 otherwise	0.144	–
County	1 for a county entity and 0 otherwise	0.059	–
School	1 for a school district and 0 otherwise	0.292	–
Town	1 for a town entity and 0 otherwise	0.168	–
Village	1 for a village entity and 0 otherwise	0.232	–
Others	1 for all remaining entity types and 0 otherwise	0.106	–
Coverage	Coverage amount in millions of dollars	35.107	78.886

Table 1 shows substantial dispersion in all three outcomes. The mean claim frequency is 0.589 claims per policyholder-year. It should not be interpreted as three claims annually; multiplying by five gives approximately 2.945 claims over the complete observation window for a continuously observed entity. Mean individual positive severity is approximately $19,515. The large standard deviations of the deductible measure and severity motivate flexible positive heavy-tailed margins.

Figure 1 displays unconditional exploratory patterns. Panel (a) suggests that policyholder-years with more claims tend to have lower deductible ratios. This pattern is consistent with selection on risk retention, but the plot alone cannot distinguish adverse selection from observed or omitted portfolio heterogeneity. Panel (b) shows that larger deductible ratios are associated with larger observed losses among positive claims. A reporting or truncation mechanism is one possible explanation because a higher deductible raises the threshold at which a loss becomes

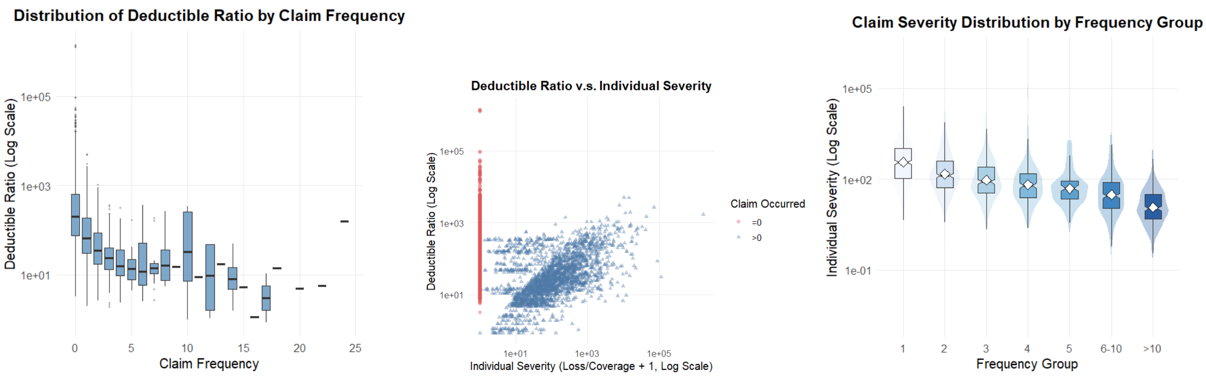

(a) Deductible ratio and claim frequency **(b)** Deductible ratio and normalized positive loss **(c)** Positive severity by frequency

Fig. 1: Exploratory relationships among the deductible ratio, claim frequency, and claim outcome. Panel (b) includes zero-claim policyholder-years for descriptive comparison; these observations do not enter the individual positive-severity likelihood.

economically relevant to report. The same pattern could also reflect entity composition, coverage design, or unobserved risk characteristics; it is therefore not treated as causal evidence of ex-post moral hazard. Panel (c) indicates a negative unconditional association between claim count and individual positive severity, supporting the use of a dependent frequency-severity specification [3–5].

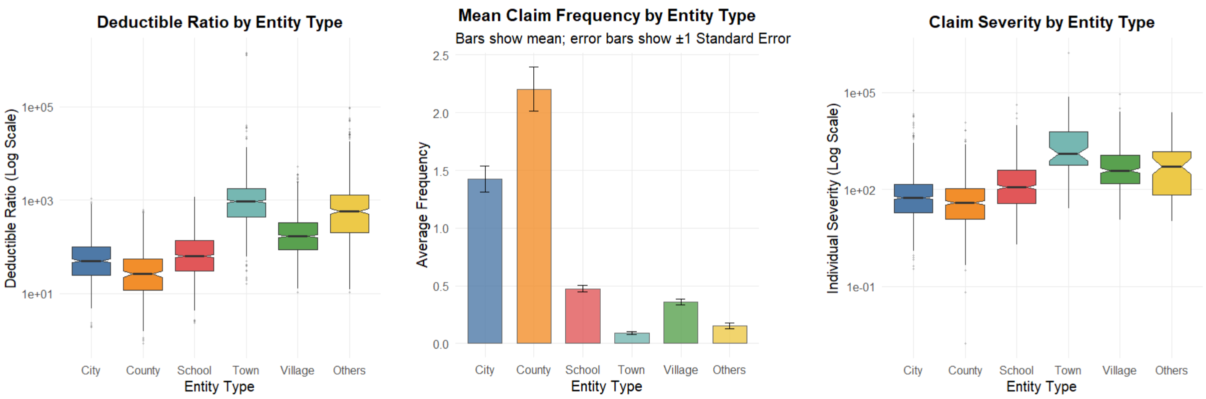

(a) Deductible ratio by entity type **(b)** Mean claim frequency by entity type **(c)** Positive claim severity by entity type

Fig. 2: Heterogeneity in deductible, claim frequency, and positive severity across entity types.

Figure 2 demonstrates that entity classification is a potentially important confounder. Towns and other entities tend to carry higher deductibles, whereas cities and counties carry lower deductibles. Counties have comparatively high claim frequencies, while towns have lower frequencies but larger positive losses. The regression models therefore include entity indicators and coverage. Any dependence remaining after this adjustment is described as residual conditional association; it may still contain omitted heterogeneity.

3.2. Marginal and copula estimates

The primary model uses the normalized deductible ratio $R = D/C$ rather than the absolute contractual deductible. Both R and individual positive severity are fitted with Generalized Beta of the Second Kind regressions, while count is fitted with a Zero-One Inflated Negative Binomial regression. The latter combines a negative-binomial component with additional masses at zero and one, reflecting the pronounced concentration at those counts [9, 15]. The "Others" category is the reference entity class, and the coverage covariate enters the frequency and severity models on the logarithmic scale.

The reported estimates are obtained by the staged procedure in Eq. (13); they are not simultaneous full-information MLEs. Standard errors are conditional observed-information standard errors from the corresponding estimation stage.

Table 2: Stage-wise marginal parameter estimates for the dependence model

Marginal models								
Deductible ratio			**Frequency**			**Positive severity**		
Parameter	Estimate	Std. error	Parameter	Estimate	Std. error	Parameter	Estimate	Std. error
Intercept	5.857	0.118	Intercept	-2.855	0.249	Intercept	7.524	0.214
City	-2.295	0.076	City	1.289	0.210	City	-0.249	0.155
County	-2.927	0.097	County	1.387	0.217	County	-0.537	0.166
School	-2.088	0.069	School	0.143	0.204	School	-0.241	0.156
Town	0.583	0.074	Town	0.827	0.358	Town	-0.144	0.202
Village	-1.066	0.071	Village	0.952	0.226	Village	-0.113	0.152
σ	0.845	0.094	log(Coverage)	0.519	0.046	log(Coverage)	0.023	0.026
α_1	1.698	0.319	ϕ	0.879	0.083	σ	0.863	0.110
α_2	1.215	0.197	**Zero-inflation model**			α_1	2.092	0.466
			Intercept	0.281	1.022	α_2	0.526	0.098
			log(Coverage)	-1.070	0.249			
			One-inflation model					
			Intercept	-2.438	0.985			
			log(Coverage)	-0.354	0.288			

Table 2 shows clear differences in risk retention across entity types relative to the “Others” category. Counties and cities have negative deductible-ratio coefficients, while towns have a positive coefficient. Coverage has a positive association with frequency but a small estimated association with individual severity. These regression coefficients describe conditional mean or scale relationships under the stated marginal parameterizations; they should not be interpreted as causal effects of entity status or coverage.

Table 3: Candidate copulas for deductible–severity and conditional frequency–severity dependence

Copula (R, Y)	Copula $(N, Y \mid R)$	NegLogLik	AIC	BIC
Frank	Rotated Gumbel 90°	31709.776	63443.551	63513.168
Frank	Gaussian	31710.970	63445.941	63515.558
Gumbel	Rotated Gumbel 90°	31714.079	63452.157	63521.774
Gumbel	Gaussian	31715.106	63454.212	63523.829
Survival Clayton	Gaussian	31716.611	63457.221	63526.838

Table 4: Candidate copulas for deductible–frequency dependence

Copula (R, N)	NegLogLik	AIC	BIC
Student-t	3453.594	6911.189	6923.851
Frank	3458.699	6919.398	6925.729
Gaussian	3458.754	6919.508	6925.839
Rotated Gumbel 270°	3461.259	6924.518	6930.849
Rotated Gumbel 90°	3462.005	6926.010	6932.341

Tables 3 and 4 report negative log-likelihood values; lower values are preferred. The AIC values are consistent with $2\,\text{NegLogLik} + 2k$. Among the evaluated candidates, a Student-t copula is selected for (R, N), a Frank copula for (R, Y), and a 90° rotated Gumbel copula for $(N, Y \mid R)$. These comparisons are conditional on the candidate set and the fitted marginal distributions.

Table 5: Estimated copula parameters and implied Kendall's τ

Component	Copula	Parameter	Estimate	Std. error	Kendall's τ
C_{RN}	Student-t	ρ	-0.2689673	0.0202909	-0.173
		ν	12.3180848	4.1522202	–
C_{RY}	Frank	θ	1.0700252	0.1235319	0.118
$C_{NY\mid R}$	Rotated Gumbel 90°	θ_{signed}	-1.5041262	0.0419661	-0.335

The Student-t estimate in Table 5 implies a moderate negative rank association between the deductible ratio and count. The estimated degrees of freedom should not be interpreted, on its own, as evidence of strong extreme co-movement; the fitted model's same-tail dependence is limited at the reported parameter values. The Frank parameter is denoted by θ, not by a correlation coefficient, and implies a weak positive rank association between the deductible ratio and observed positive severity. For the rotated Gumbel, the negative signed parameter follows the software convention for a 90° rotation; the corresponding unrotated Gumbel magnitude is 1.5041 and the implied rank association is negative.

The signs are compatible with a lower deductible ratio among observations with higher claim frequency, larger observed positive losses at higher deductible ratios, and smaller positive losses at higher claim counts after conditioning on the deductible ratio. These are conditional associations. The first may be consistent with selection on latent risk, while the second may be consistent with reporting thresholds or selective reporting. Neither identifies adverse selection or moral hazard without additional instruments, structural restrictions, or a longitudinal behavioral design.

As a support-sensitivity check, the deductible was also represented by its ordered contractual level rather than by the continuous ratio R. The ordinal rerun retained the same three selected copula families Student-t for (R, N), Frank for (R, Y), and 90° rotated Gumbel for $(N, Y \mid R)$ and the same dependence directions. For quantitative reference, the corresponding primary-model Kendall's τ values are -0.173, 0.118, and -0.335, respectively. The holdout evaluation used for the primary continuous specification contains 1,038 policyholder-year observations; the full model wins on 64.55% of frequency log-score comparisons and 71.09% of severity comparisons, with a 22.54-point global Insurance Gini gain. These latter predictive figures are reported for the primary continuous specification, not as separate ordinal-model estimates. The ordinal exercise is therefore interpreted as a robustness check on the support assumption rather than as a second preferred fitted model.

3.3. Out-of-sample predictive validation

The temporal validation compares the full dependence model with the exogenous-deductible benchmark defined in Eq. (11). For frequency, the predictive distributions are $p_N(n \mid \mathbf{x})$ and $p_{N|R}(n \mid r; \mathbf{x})$. For severity, the corresponding distributions are $f_{Y|N}(y \mid n; \mathbf{x})$ and $f_{Y|RN}(y \mid r, n; \mathbf{x})$. The 2010 observations are not used for parameter estimation.

Table 6: Out-of-sample log-score comparison on the 2010 temporal holdout

Outcome	Full-model win rate	Exogenous-deductible total score	Full-dependence total score
Claim frequency	64.55%	-1049.774	-1022.059
Claim severity	71.09%	-8366.561	-8214.455

Table 6 reports total, rather than average, logarithmic scores. A higher score is preferred. The full model assigns a higher observation-level log score than the benchmark for 64.55% of count observations and 71.09% of severity observations (670 and 738 of the 1,038 holdout observations, respectively). These percentages are empirical win rates defined in Eq. (16). No formal sampling uncertainty or hypothesis test is attached to these proportions in this study; they are descriptive summaries of the fixed 2010 temporal holdout, not classification accuracies and not probabilities that the model is correct. The higher total scores indicate improved predictive density assignment on the holdout, but the result should be understood as temporal validation on entities that may also appear in the estimation years.

4. Managerial Relevance and Insurance Gini Evaluation

The empirical results have two distinct operational interpretations. First, the deductible ratio can provide incremental information for one-year-ahead risk ranking after observed rating variables have been considered. Second, this incremental information does not establish that changing a policyholder's deductible would causally change claim behavior. Accordingly, the discussion below focuses on conditional prediction and discrimination rather than guaranteed underwriting profit.

4.1. Conditional risk relativities

Let S_i denote aggregate loss, D_i the contractual deductible, and $g(S_i, D_i)$ the insurer's payment function. For covariates $\mathbf{x}_i$, define the full-model conditional expected payment and its exogenous-deductible counterpart as

$$m_{\text{full}}(r, \mathbf{x}_i) = E[g(S_i, D_i) \mid R_i = r, \mathbf{x}_i], \qquad m_{\text{exo}D}(\mathbf{x}_i) = E[g(S_i, D_i) \mid \mathbf{x}_i]. \tag{17}$$

For a reference ratio r_0, a conditional risk relativity may be written as

$$\text{rel}(r; r_0, \mathbf{x}_i) = \frac{m_{\text{full}}(r, \mathbf{x}_i)}{m_{\text{full}}(r_0, \mathbf{x}_i)}. \tag{18}$$

Equation (18) is a predictive relativity. It measures how the fitted conditional loss cost changes with the observed deductible ratio within the model; it is not a treatment effect of assigning a different deductible.

4.2. Champion–challenger Insurance Gini

The 2010 holdout is also evaluated using the ordered-Lorenz Insurance Gini of Frees, Meyers, and Cummings [16]. Let $p_i^B > 0$ be the base-model score, $p_i^C > 0$ the challenger score, and $s_i \geq 0$ the realized loss measure. Define $Q_i = p_i^C / p_i^B$ and order observations so that $Q_{(1)} \leq \cdots \leq Q_{(n)}$. The cumulative base-score and loss shares are

$$P_k = \frac{\sum_{h=1}^{k} p_{(h)}^B}{\sum_{i=1}^{n} p_i^B}, \qquad L_k = \frac{\sum_{h=1}^{k} s_{(h)}}{\sum_{i=1}^{n} s_i}, \qquad P_0 = L_0 = 0. \tag{19}$$

Using the sign convention in this study, the pairwise statistic in Gini points is evaluated by trapezoidal integration:

$$G(B, C) = 200 \sum_{k=1}^{n} \left[\frac{(P_{k-1} - L_{k-1}) + (P_k - L_k)}{2} \right] (P_k - P_{k-1}). \tag{20}$$

A positive $G(B, C)$ means that the column/challenger score provides better ordering than the row/base score under this convention. Because the base weights and ordering change when the models are reversed, $G(B, C)$ need not equal $-G(C, B)$.

In the implementation, the realized loss measure is $s_i = \log(1 + \text{Claim}_i)$, and all competing risk scores are evaluated on the corresponding transformed scale. Consequently, the Gini values quantify relative discrimination on that scale; they should not be read as a dollar profit margin. A constant score represents an uninformative or tied benchmark, not a mathematical "maximum possible loss." An oracle score based on realized outcomes is useful only as an unattainable reference bound.

Table 7: Pairwise Insurance Gini comparisons in points; standard errors are in parentheses

Base \ Challenger	Pooled	Exogenous deductible	Full dependence	Experience
Pooled	–	36.5633	37.9323	35.3548
		(2.0120)	(2.0081)	(2.0339)
Exogenous deductible	-27.4586	–	22.5442	9.5223
	(2.0520)		(2.1214)	(2.2387)
Full dependence	-27.5477	-20.1313	–	-9.4131
	(2.0580)	(2.1233)		(2.1905)
Experience	-25.1190	-6.6048	12.5025	–
	(2.0765)	(2.2356)	(2.1806)	

Four scores are compared: a pooled score that ignores observable heterogeneity, an exogenous-deductible score that removes R–claim links while retaining frequency–severity dependence, the full dependence score, and the Fund's experience or contract score.

Table 7 shows that the full dependence score improves on the exogenous-deductible score by 22.5442 Gini points (SE $= 2.1214$), corresponding to a nominal 95% interval of approximately $[18.39, 26.70]$. It also provides better ordering than the experience score when the latter is used as the base. These comparisons support an incremental discrimination benefit from conditioning on the deductible ratio. They do not establish causal asymmetric information or guarantee realized underwriting profit.

4.3. Deductible-stratified evaluation

To examine heterogeneity in discrimination, the 1,038 holdout entities are partitioned using prespecified nominal deductible bands. The cutoffs are held fixed across models so that the comparison is not selected using model performance.

Table 8: Exogenous-deductible versus full-dependence Insurance Gini by contractual deductible stratum (US dollars)

Stratum	Contractual deductible	n	$G(\text{exo}D, \text{full})$	Std. error	Nominal 95% CI	$G(\text{full}, \text{exo}D)$
1	$D \leq 500$	447	25.03	3.385	$[18.39, 31.67]$	-21.87
2	$500 < D \leq 1{,}000$	220	21.94	4.232	$[13.65, 30.23]$	-19.68
3	$1{,}000 < D \leq 2{,}500$	133	16.13	6.207	$[3.96, 28.30]$	-14.53
4	$2{,}500 < D \leq 5{,}000$	143	25.88	5.468	$[15.16, 36.60]$	-24.12
5	$D > 5{,}000$	95	14.61	7.880	$[-0.83, 30.05]$	-13.18

All five point estimates in Table 8 favor the full dependence score. The evidence is strongest in strata 1, 2, and 4, while the smaller third stratum has a wider interval. For stratum 5, the nominal 95% interval includes zero and the corresponding z statistic is approximately 1.85. The highest-deductible result is therefore described as positive but imprecise, rather than statistically significant at the 5% level. The intervals are pointwise and no claim of simultaneous significance across all five groups is made.

Figure 3 is consistent with the point estimates in Table 8. The stratified analysis indicates that the ranking gain is not concentrated exclusively in the lowest deductible group. At the same time, the widening uncertainty as the stratum size decreases cautions against a universal profitability claim. The operational implication is that the deductible ratio can be evaluated as an additional segmentation variable, subject to calibration, governance, and stability checks before implementation.

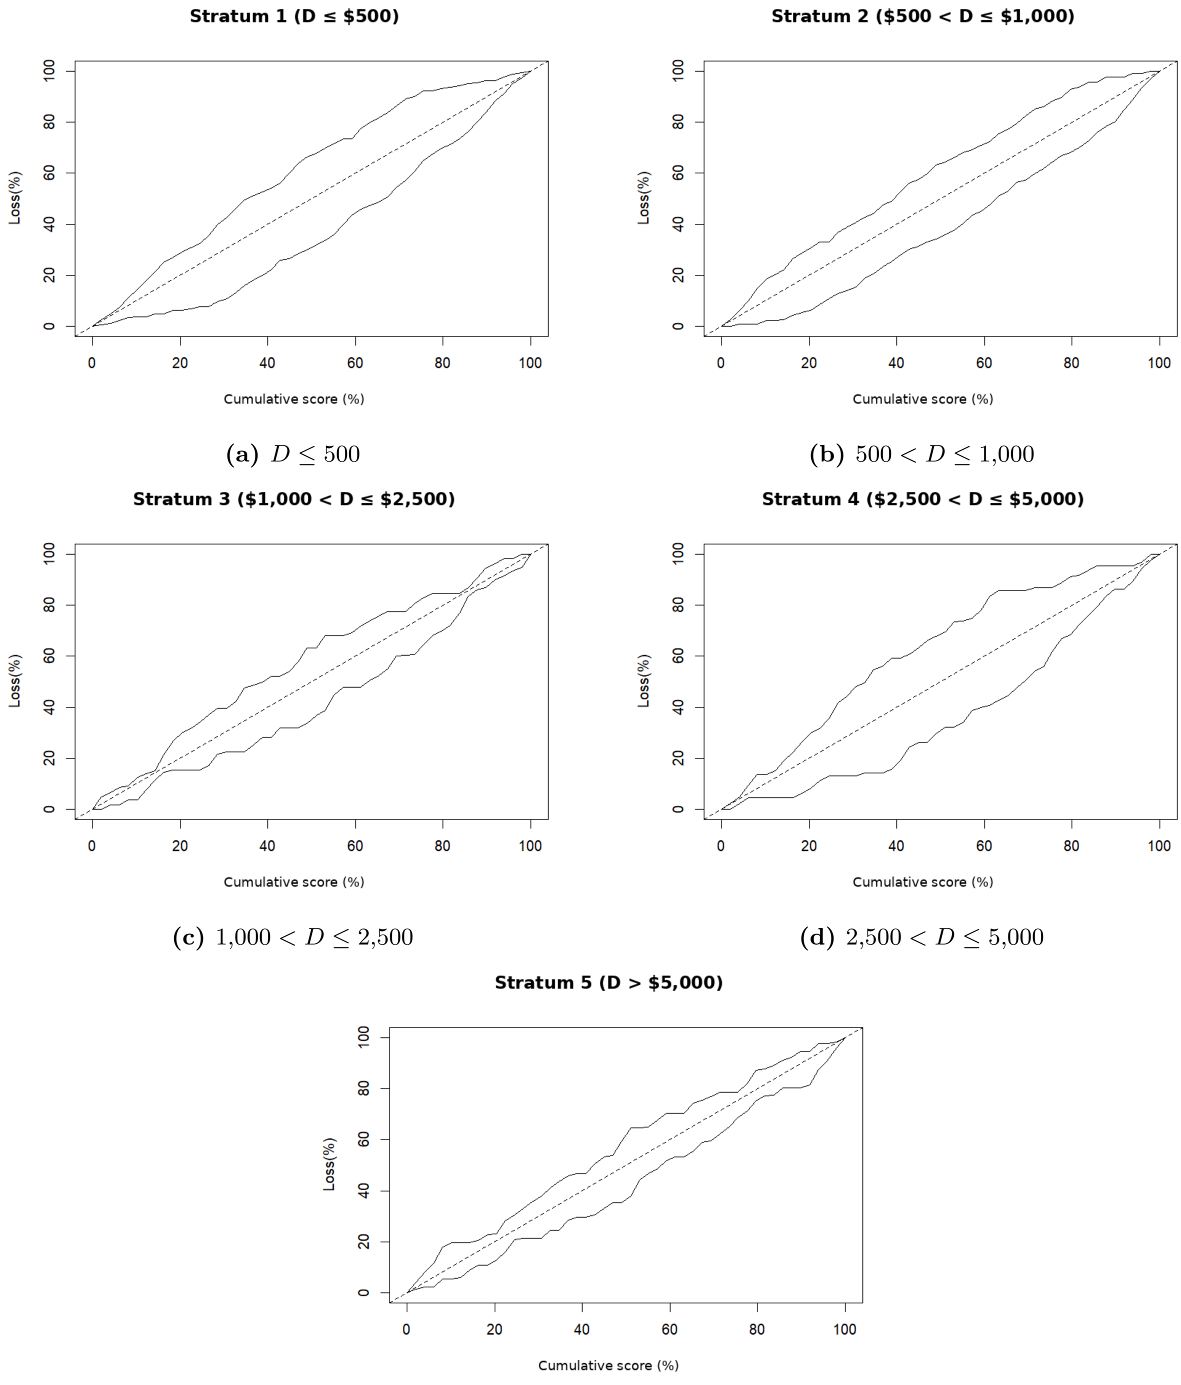

(a) $D \leq 500$ **(b)** $500 < D \leq 1{,}000$

(c) $1{,}000 < D \leq 2{,}500$ **(d)** $2{,}500 < D \leq 5{,}000$

(e) $D > 5{,}000$

Legend: upper solid curve = full-dependence score; lower solid curve = exogenous-deductible score; - - - = equality line.

Fig. 3: Ordered-Lorenz comparisons of the exogenous-deductible and full-dependence scores across deductible strata. The horizontal axis is the cumulative percentage of the model score (premium/risk score) after ordering observations by the challenger-to-base score ratio; the vertical axis is the corresponding cumulative percentage of realized loss. The constant-score reference appearing in the panels is interpreted as an uninformative benchmark, not as a mathematical maximum-loss bound.

5. Conclusion

This study revisited the copula regression framework of Shi and Lee [9] and extended its empirical evaluation across deductible strata. The analysis distinguished the contractual deductible from the normalized continuous ratio used in the primary model, retained a discrete likelihood contribution for claim count, and stated explicitly how the random number of individual severities enters the policyholder-year likelihood. It also separated the exogenous-deductible benchmark, which retains frequency–severity dependence, from a total-independence model.

The fitted C-vine produced negative deductible–frequency dependence, positive deductible–observed-severity dependence, and negative conditional frequency–severity dependence. These estimates show residual conditional associations after observed rating variables have been included. The negative deductible–frequency association is compatible with selection on latent risk, while the positive deductible–observed-severity association is compatible with reporting thresholds or selective reporting. Neither result constitutes causal identification of adverse selection or ex-post moral hazard.

On the 2010 temporal holdout, the full dependence model achieved a higher observation-level logarithmic score than the exogenous-deductible benchmark for 64.55% of frequency observations and 71.09% of severity observations. The full score also improved the global Insurance Gini comparison by 22.54 points. The deductible-stratified point estimates were positive in all five groups, ranging from 14.61 to 25.88 points. However, the interval for the highest-deductible group included zero, so the evidence does not support a claim of statistically significant improvement in every stratum. The Gini analysis measures relative ordering on the evaluation scale and should not be interpreted as a guaranteed dollar profit.

Several limitations remain. The continuous GB2 model for $R = D/C$ is an approximation to a contract choice generated from a finite deductible menu, although the ordinal sensitivity check preserves the qualitative conclusions. Point estimation uses a pooled working likelihood for a short panel, and the conditional stage-wise standard errors do not fully propagate first-stage and serial-dependence uncertainty. The temporal holdout contains entities that may have appeared during training, and the observational design cannot isolate causal behavioral mechanisms. Observed severity may also be affected by reporting or recording thresholds. Finally, discrimination should be supplemented by calibration and operational monitoring before the score is used for pricing decisions.

Future research should model the repeated observations directly through random effects or a dynamic panel copula, propagate uncertainty across all estimation stages, and evaluate deductible support using a fully discrete or ordinal likelihood. The panel framework of Shi and Zhang [10] provides a natural reference for that extension. Additional applications to independent portfolios and prospective policy periods would help determine whether the stratified gains generalize beyond the LGPIF data.

CRediT Authorship Contribution Statement

Dwi Mifta Mahanani: Supervision and Investigation. **Feby Indriana Yusuf:** Conceptualization, Methodology, Writing – Original Draft Preparation, Project Administration, Funding Acquisition. **Dzaki Ferlian Nugroho:** Data Curation, Formal Analysis, Software, Validation, Visualization, Writing – Original Draft, Writing – Review & Editing. **Tuti Sariningsih Budi Utami:** Validation and Investigation.

Declaration of Generative AI and AI-assisted technologies

During the preparation of this work, Generative AI tools (specifically GeminiAI) were used solely for language refinement and grammar editing purposes. No part of the analysis, interpretation, or core content was generated by AI. All substantive results and discussions were manually developed by the authors.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding and Acknowledgments

This research was funded by the Faculty of Science, Technology, and Mathematics, Universitas Brawijaya, through the FSTeM Research Grant Scheme B. The authors gratefully acknowledge the Faculty of Science, Technology, and Mathematics, Universitas Brawijaya, for its financial support in conducting this research and preparing this publication.

Data and Code Availability

The LGPIF data and the computational foundation adapted in this study are documented by Shi and Lee [9], available through DOI 10.1080/01621459.2022.2040519, and in the public repository at `https://github.com/user-pshi/CompoundRegression`. The present manuscript's deductible-stratified evaluation and reporting revisions were developed by the current authors. Reproducibility materials for these extensions are available from the corresponding author upon reasonable request. Any reused or adapted code should retain attribution to the original repository and article.

References

[1] Pierre-André Chiappori and Bernard Salanié. "Testing for Asymmetric Information in Insurance Markets". In: *Journal of Political Economy* 108.1 (2000), pp. 56–78. DOI: 10.1086/262111.

[2] Alma Cohen and Peter Siegelman. "Testing for Adverse Selection in Insurance Markets". In: *Journal of Risk and Insurance* 77.1 (2010), pp. 39–84. DOI: 10.1111/j.1539-6975.2009.01337.x.

[3] Edward W. Frees, Gee Lee, and Lu Yang. "Multivariate Frequency–Severity Regression Models in Insurance". In: *Risks* 4.1 (2016), p. 4. DOI: 10.3390/risks4010004.

[4] Gee Y. Lee and Peng Shi. "A Dependent Frequency–Severity Approach to Modeling Longitudinal Insurance Claims". In: *Insurance: Mathematics and Economics* 87 (2019), pp. 115–129. DOI: 10.1016/j.insmatheco.2019.04.004.

[5] Peng Shi and Zifeng Zhao. "Regression for Copula-Linked Compound Distributions with Applications in Modeling Aggregate Insurance Claims". In: *The Annals of Applied Statistics* 14.1 (2020), pp. 357–380. DOI: 10.1214/19-AOAS1299.

[6] Kjersti Aas, Claudia Czado, Arnoldo Frigessi, and Henrik Bakken. "Pair-Copula Constructions of Multiple Dependence". In: *Insurance: Mathematics and Economics* 44.2 (2009), pp. 182–198. DOI: 10.1016/j.insmatheco.2007.02.001.

[7] Peng Shi and Lu Yang. "Pair Copula Constructions for Insurance Experience Rating". In: *Journal of the American Statistical Association* 113.521 (2018), pp. 122–133. DOI: 10.1080/01621459.2017.1330692.

[8] Lu Yang, Edward W. Frees, and Zhengjun Zhang. "Nonparametric Estimation of Copula Regression Models with Discrete Outcomes". In: *Journal of the American Statistical Association* 115.530 (2020), pp. 707–720. DOI: 10.1080/01621459.2018.1546586.

[9] Peng Shi and Gee Y. Lee. "Copula Regression for Compound Distributions with Endogenous Covariates with Applications in Insurance Deductible Pricing". In: *Journal of the American Statistical Association* 117.539 (2022), pp. 1094–1109. DOI: 10.1080/01621459.2022.2040519.

[10] Peng Shi and Wei Zhang. "A Copula-Based Panel Data Model for Pricing Insurance Contracts with Endogenous Deductibles". In: *The Econometrics Journal* (2026), utag018. DOI: 10.1093/ectj/utag018.

[11] Christian Genest and Johanna Nešlehová. "A Primer on Copulas for Count Data". In: *ASTIN Bulletin* 37.2 (2007), pp. 475–515. DOI: 10.2143/AST.37.2.2024077.

[12] Whitney K. Newey and Daniel McFadden. "Large Sample Estimation and Hypothesis Testing". In: *Handbook of Econometrics*. Ed. by Robert F. Engle and Daniel L. McFadden. Vol. 4. Elsevier, 1994, pp. 2111–2245. DOI: 10.1016/S1573-4412(05)80006-4.

[13] Tilmann Gneiting and Adrian E. Raftery. "Strictly Proper Scoring Rules, Prediction, and Estimation". In: *Journal of the American Statistical Association* 102.477 (2007), pp. 359–378. DOI: 10.1198/016214506000001437.

[14] Claudia Czado, Tilmann Gneiting, and Leonhard Held. "Predictive Model Assessment for Count Data". In: *Biometrics* 65.4 (2009), pp. 1254–1261. DOI: 10.1111/j.1541-0420.2009.01191.x.

[15] Jean-Philippe Boucher, Michel Denuit, and Montserrat Guillén. "Number of Accidents or Number of Claims? An Approach with Zero-Inflated Poisson Models for Panel Data". In: *Journal of Risk and Insurance* 76.4 (2009), pp. 821–846. DOI: 10.1111/j.1539-6975.2009.01321.x.

[16] Edward W. Frees, Glenn Meyers, and A. David Cummings. "Summarizing Insurance Scores Using a Gini Index". In: *Journal of the American Statistical Association* 106.495 (2011), pp. 1085–1098. DOI: 10.1198/jasa.2011.tm10506.