

SCOR-UQAM NEWSLETTER #5 (October 2025-March 2026)

Arthur Charpentier (UQAM; charpentier.arthur@uqam.ca)

After a busy summer, and as the project enters its final year, Arthur Charpentier was invited to spend a few months in Japan, at 京都大学, Kyoto University (or [KyotoU](#)), in Kyoto, 京大. In particular, he gave the course “Five Lectures on Fairness and Discrimination in Insurance: An Actuarial Perspective” to graduate students at RIMS (*Research Institute for Mathematical Sciences*) and Kyoto University. Meanwhile, this fall and winter, Agathe Fernandes Machado and Olivier Côté, the two PhD students involved in this project, gave many talks during their third year of doctoral studies (we will return to that from page 12). Congratulations to them.

Among the recent work presented in this newsletter, we revisit [Linear Risk Sharing on Networks](#), where we explain how to share risk on graphs using doubly stochastic linear operators. [Decomposing Probabilistic Scores: Reliability, Information Loss and Uncertainty](#), which decomposes probabilistic scores into calibration, information loss, and residual uncertainty. [Perceived Fairness in Networks](#) proposes a way to model fairness as it is locally perceived in a network, when individuals compare themselves to their peers. [A Three-Pillar Governance Framework for Fair Insurance Pricing](#) develops a three-pillar governance framework for fair insurance pricing (balancing solidarity, causality, and actuarial fairness). [Decomposing Direct and Indirect Biases in Linear Models under Demographic Parity Constraint](#) proposes a method to distinguish direct and indirect bias in linear models under demographic parity constraints. [Fair regression under localized demographic parity constraints](#) enforces demographic parity in regression through quantile constraints. In [Sequential Transport for Causal Mediation Analysis](#), we analyze causal mediation through sequential transport in order to better attribute disparities. We also return to data and inclusive insurance in [Criminal Records, the Right to be Forgotten, and Inclusive Insurance](#), focusing on the specific case of criminal records, the right to be forgotten, and information gov-

ernance. In [Delving into Deep Smoothed Bootstrap: Application in Imbalanced Regression](#), we improve learning in imbalanced regression using VAEs and *smoothed bootstrap*. In [KurtHGR: A Neural Maximal Correlation for Tabular Datasets](#), we estimate maximal correlation (HGR) with neural methods for tabular data. [Improving Procedural and Substantive Fairness in Conformal Prediction](#) attempts to connect procedural and substantive fairness in conformal prediction, with *LLM-in-the-loop* evaluation. And finally, two general-audience articles in [Risques, The Impossible Right to Make Mistakes, the Impossible Right to Be Forgotten?](#) and, more recently, [What If We Finally Invested in Prevention?](#).



Associated People

Students

With all our travel, there will be no intern next summer. However, even though the project will come to an end this fall, we will recruit for Summer 2027, and applications will need to be submitted early (starting in September), via the platform <https://globalink.mitacs.ca/>.

Also, this winter, Arthur served on the jury for the first edition of the “[My Apprenticeship/Internship in 180 Seconds](#)” competition, organized by the **Association des Masters en Actuariat** in France. Very nice presentations, congratulations to all participants, 🎉

📍 Where are they now?

[Raphaël Dalbarade](#) , former intern of Arthur and Laurence Barry, and [Bertille Tierny](#) , former intern of François Hu who was also supervised by Arthur, defended their actuarial theses in France at ENSAE Institut Polytechnique, respectively on [Study and Measurement of the Impact of Insurers’ Withdrawal from Areas at High Risk of Natural Catastrophes](#) (for Raphaël) and [Fairness, Audit, and Regulatory Compliance of Insurance Models: An Analytical Framework for Application to the Colorado Regulation](#) (for Bertille), 🎉

See <https://freakonometrics.hypotheses.org/>

Lucas Offroy [in](#), a former intern of Arthur and Marouane Il-Idrissi, and **Agathe Fernandes Machado** spoke to students at the actuarial school of the University of Western Brittany (EURIA) about the explainability of predictive models.



Sang Huynh (alumni: MSc University of Lyon and National Economics University of Hanoi, Vietnam) is advancing his work on fairness and multivariate optimal transport, under the supervision of **Vương Văn Yên** (professor in Hà Nội, Vietnam) and Arthur Charpentier, [in](#)

Sang Huynh was awarded the **2025–2026 CAS Trust Scholarship** from the Casualty Actuarial Society 🧐

Iryna Voitsitska [in](#), a former intern, saw her work **ElliCE: Efficient and Provably Robust Algorithmic Recourse via the Rashomon Sets** (with Bohdan Turbal, Princeton University, and Lesia Semenova, Rutgers University) presented this fall at NeurIPS'25, 🧐

Cassandra Mussard [in](#), a former intern, began a PhD this fall at the Centre d'Études Spatiales de la BIOSphère (Cesbio) in France, on deep learning techniques, 🧐

Suzie Grondin [in](#), a former intern, began a PhD this fall at Sorbonne Université (LIP6), on the interpretability of predictions. We should cross paths again in the coming years... 🧐



From left to right: Suzie Grondin, François Hu, Agathe Fernandes Machado, Bertille Tierney, and Raphaël Dalbarade, in Paris.

Thanks to the team's alumni for keeping us regularly informed, and for sending photos from time to time!

PhD Students

A “new” addition to the team...

Raphaël Dalbarade (alumni: MSc ENSAE, BSc Université Paris Saclay, France) began a PhD under the supervision of **Caroline Hillairet** (ENSAE Institut Polytechnique) and Arthur Charpentier, [in](#)



The topic of the thesis is “**Insurability of Climate Risk between Segmentation and Pooling: Game Theory, Competition, Equilibria, and Public Policy.**”

Agathe Fernandes Machado was a visiting researcher at Aix-Marseille University for four months. She gave talks in several seminars (France, the Netherlands, and Scotland). **Olivier Côté** also spent time in Europe (Switzerland, Belgium, and Sweden).

📍 Where are they now?

Samuel Stocksieker (Newsletter #2 ↪️) received the “2025 Young Doctors Prize” for his doctoral thesis, “*Contribution of machine learning in modeling rare values and imbalanced data: applications in insurance,*” this winter. Congratulations Sam 🧐



Postdoctoral Fellows

Arsène Brice Zotsa Ngoufac (Newsletter #3 ↪️) was awarded a 3-year fellowship at **Vanderbilt University**, Nashville, Tennessee, ranked among the top fifteen* universities in the United States. He left the team at the beginning of spring, 🧐

For **Marouane Il-Idrissi** (Newsletter #3 ↪️), more information in the next newsletter.

Miscellaneous



Marie-Pier Côté, professor at Université Laval in Quebec City, [in](#), with whom we have collaborated for several years, received an **NSERC Alliance** grant, “**Refinements of general insurance pricing models**,” with Thierry Duchesne, [👤](#)

Caroline Hillairet, professor at ENSAE, is Raphaël’s PhD advisor, with Arthur as co-supervisor. They are organizing a conference in Mérida, Yucatan (Mexico) next fall (just after *día de muertos*).



The conference **Mathematical Advances on Emerging Risks** will be organized with Thomas Peyrat (ENSAE and Institut de Mathématiques de Toulouse), Dylan Posamaï (ETH Zürich), Anthony Réveillac (INSA Toulouse), Daniel Hernández (CIMAT), Ricardo Guerrero Joshué Heli (ETH Zürich), and Renata Gomes Alcoforado (Universidade Federal de Pernambuco) [🌐](#). This conference continues the partnerships developed with Mexico over the past few months, which we mentioned in [Newsletter #3](#) [👉](#). We will return to this event in the next newsletter.

Caroline was also involved in organizing the **Printemps de l’Assurance** (5th edition), which took place in late March in Paris.

Arthur Charpentier and **Bertrand Villeneuve**, professor at Paris Dauphine University, [in](#), are involved in a project supported by the ACPR Chair (Autorité de contrôle prudentiel et de résolution) of the Banque de France, to work on insurance supply and demand [👤](#)



Finally, taking advantage of his invitation to Kyoto, Arthur Charpentier took part in activities with local groups, in particular at 京都大学, at the **Disaster Prevention Research Institute**, in Kyoto.

Recent Work

As mentioned in the introduction, we completed several projects on explainability and attribution, score calibration, and risk-sharing architectures, along with a series of papers on fairness in pricing and its links with causality, optimization, and distributed learning. We also worked on data and inclusive insurance, multivariate optimal transport, estimation under measurement error, social welfare issues, LLM (*Large Language Models*)-based evaluation, and climate risk. But first, a look back at the work published this fall and winter.

Publications

KNN and K-Means in Gini Prametric Spaces by Cassandra Mussard, Arthur Charpentier and Stéphane Mussard ([Newsletter #3](#) [👉](#)) was published in *Frontiers in Artificial Intelligence and Applications*,

[doi:10.3233/FAIA251085](#),



The papers **KurtHGR: Neural Maximal Correlation for Tabular Datasets** and **Disentangled Deep Smoothed Bootstrap for Fair Imbalanced Regression** by Samuel Stocksieker, Denys Pommeret and Arthur Charpentier were published in *Procedia Computer Science*

[doi:10.1016/j.procs.2025.09.519](#),

[doi:10.1016/j.procs.2025.09.338](#),



Decomposing Direct and Indirect Biases in Linear Models Under Demographic Parity Constraint by Bertille Tierny, Arthur Charpentier and François Hu (presented on page 9) was published in *Proceedings of the AAAI Conference on Artificial Intelligence*,

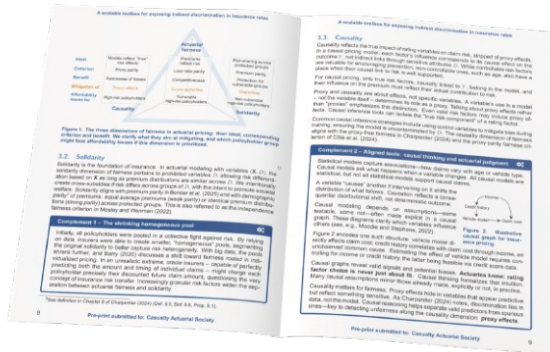
[doi:10.1609/aaai.v40i31.39793](#),



The study **A Scalable toolbox for exposing indirect discrimination in insurance rates**, by Olivier Côté, Marie-Pier Côté and Arthur Charpentier, was finally published in the *Casualty Actuarial Society* collection,



casact.org/sites/default/files/2026-03/...



In the *Revue Risques*, publication of **The Impossible Right to Make Mistakes, the Impossible Right to Be Forgotten?**



Arthur Charpentier, Olivier Côté and Marie-Pier Côté also published, in the PARI Chair working paper series, **A Three-Pillar Governance Framework for Fair Insurance Pricing WP #37**, following **WP #32** with Xavier Vamparys, and **WP #26** with Laurence Barry, in the same series.

For the cover illustration, after a series of Flemish painters (Petrus Christus and Jan Steen), this time we have 勝川春潮, Katsukawa Shunchō, on the cover (the original being stored in a drawer at the MET in New York).




Arthur Charpentier was mentioned in the conclusion of the **2025 Annual Report on the State of France**, written by the **Conseil Économique, Social et Environnemental (CESE)** in France and published this fall.



“The worsening of extreme events linked to climate change requires a strengthening of our insurance protection system, while maintaining a balance that is sustainable and acceptable for all types of stakeholders (individuals and firms / insurers / reinsurers / public authorities — CCR, Barnier Fund). In addition, a considerable step change in prevention efforts is needed. Climate disruption tends to deepen inequalities and calls for a stronger “social contract of risk, a new common good, neither a pure financial product nor a simple common pool, but an essential infrastructure through which society — and not only the insurer — consciously chooses the share of the climate bill it is willing to bear,” as economist Arthur Charpentier writes.

Arthur Charpentier wrote two short articles in **Pépites mathématiques et scientifiques**, edited by Normand Baillargeon and Robert Durocher, on the **Simpson paradox** and the **Condorcet paradox** (mentioned in **Newsletter #4**).



We mentioned it in the previous newsletter, but the **freakonometrics** blog resumed its monitoring and reflection activity, with several posts, including **Is Insurance a Zero-Sum**

Game?, “Derisking,” Word of the Year, Insuring AI. New Risks? New Models?, When Insurance Unravels: The Silent Crisis of Compensation, Insurance, Real Estate, and the Manufacture of Ignorance, The State Will Pay for It!, and Correlation and Causality, Circumstance and Context (several were also translated into English).

An article, **Insurers and AI: A Risk That Has Become Too Large to Cover**, published by Novethic, revisits blog posts published by Arthur, *Insurers and AI, a Systemic Risk and Insuring AI. New Risks? New Models?*.

Hubert Guillaud, on his blog *Dans Les Algorithmes*, also discussed one of these posts, on AI insurance, *better understanding the effectiveness of AI risk mitigation measures*

“On his blog, the applied mathematics professor Arthur Charpentier revisits the two reports of the International AI Safety Report. The first explains how technical changes shift the frontier of possible uses, and therefore of plausible harms. The second explains how the sector is trying to build safety (through layers, procedures, and audits) while explicitly acknowledging that the effectiveness of these measures remains uncertain and context-dependent. [...]

then *The Hidden Cost of AI Is Verification*

“[These] “are metrics that are quite easy to surface in a dashboard, and also easy to present to senior management. But they are often very poor measures of the real quality of the work delivered,” Arthur Charpentier reminds us on his blog. [...] *Part of today’s optimism may therefore come from a very simple confusion: we confuse what is easy to count with what is actually important.* [...]

Let us now return to a few recent developments.

Linear Risk Sharing on Networks

by Arthur Charpentier and Philipp Ratz, [doi:10.48550/arXiv.2509.21411](https://doi.org/10.48550/arXiv.2509.21411).

“Over the last decade, alternatives to traditional insurance and banking have emerged, relying on network structures to redistribute risk through peer-to-peer insurance, reciprocal contracts, and decentralized finance. We propose a linear risk sharing (LRS) framework, in which random losses are reallocated by non-negative linear operators, ensuring budget balance, fairness, and diversification under conditions linked to stochastic matrices. We show that variance reduction and majorization arise naturally from doubly stochastic mixtures, then extend the analysis to networks in order to study the effect of topology and the trade-off between self-retention and diversification.

Formally, if X_i is the (random) loss of agent i , and $\mathbf{X} = (X_1, \dots, X_n)$, linear risk sharing means that there exists an $n \times n$ matrix M , with positive entries, such that after sharing, agent i faces ξ_i , where $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)$ satisfies $\boldsymbol{\xi} = M\mathbf{X}$. We have genuine sharing if the total is preserved (the balance property), i.e. $M^T \mathbf{1}_n = \mathbf{1}_n$. We have a fairness principle if each ξ_i is a convex function of the X_j , which translates into $M\mathbf{1}_n = \mathbf{1}_n$. Thus M must belong to the class of doubly stochastic matrices. In this paper, we focus on the case where M is induced by a graph on the n insureds, with Sinkhorn–Knopp normalization. More precisely, we use an average between this matrix M and the identity matrix $\mathbb{I}_{n \times n}$.

Decomposing Probabilistic Scores: Reliability, Information Loss and Uncertainty

by Agathe Fernandes Machado and Arthur Charpentier, is available at [doi:10.48550/arXiv.2603.15232](https://doi.org/10.48550/arXiv.2603.15232)

First, we show that with a loss ℓ , the average loss $[\ell(S, Y)]$ can be decomposed into a sum of *non-negative* terms that correspond to sources of error at different information levels. Let $C(s) = \mathbb{P}(Y = 1 \mid S = s)$ (the calibration function), $Q(x) = \mathbb{P}(Y = 1 \mid \mathbf{X} = x)$ (Bayes at the X level), and $\Pi(z) = \mathbb{P}(Y = 1 \mid Z = z)$ (the “true chance” at the most informative level Z). Writing $d_\ell(\cdot, \cdot)$ for the divergence associated with ℓ (KL for log-loss, squared distance

for the Brier score), we obtain schematically:

$$[\ell(S, Y)] = \mathbb{E}[H_\ell(\Pi)] + \mathbb{E}[d_\ell(S, C)] \\ + \mathbb{E}[d_\ell(C, Q)] + \mathbb{E}[d_\ell(Q, \Pi)],$$



where $H_\ell(\Pi)$ is the Bayes risk (the “entropy” associated with ℓ) at the Z level, corresponding to fundamental uncertainty. The three divergence terms can be interpreted as follows:

- **Miscalibration / reliability** $[d_\ell(S, C)]$: the gap between the announced score S and the true conditional probability given the score, $C(S)$; this is the only term that recalibration can reduce.
- **Grouping / information loss** $X \rightarrow S$, $[d_\ell(C, Q)]$: even after perfect recalibration, the compression $X \mapsto S$ groups together heterogeneous situations; this term decreases only if more information is retained in the score.
- **Latent heterogeneity** $[d_\ell(Q, \Pi)]$: even with X fixed, the “true chance” may vary (latent variables / non-stationarity), so $Q(X)$ remains an average of $\Pi(Z)$.

Thus, the theorem clearly separates what recalibration can fix (reliability) from what requires changing the score itself (grouping) or enriching the available information (latent heterogeneity). In the case of sequential ensemble approaches (boosting), a succession of decompositions can be written, with a martingale property.

“ Calibration is a conditional property that depends on the information retained by a predictor. We establish decomposition identities for arbitrary proper losses that make this dependence explicit. At any information level \mathcal{A} , expected loss decomposes into a reliability term and a conditional entropy term. For two nested levels $\mathcal{A} \subseteq \mathcal{B}$, a chain decomposition measures the information gain. In classification, with variables X and score $S = s(X)$, we obtain a three-term identity: calibration error, information loss from X to S , and irreducible uncertainty. This framework makes it possible to analyze post hoc calibration, aggregation of calibrated models, and boosting, with explicit forms for the Brier score and logarithmic loss.

A Three-Pillar Governance Framework for Fair Insurance Pricing


by Arthur Charpentier, Olivier Côté, and Marie-Pier Côté, is available online in its French version in the [PARI Chair Working Papers](#). This work conceptualizes several ideas presented in [A Scalable Toolbox for Exposing Indirect Discrimination in Insurance Rates](#), mentioned in [Newsletter #4](#) , and since published by the Casualty Actuarial Society in its *Research Paper Series on Race and Insurance Pricing*, .

“ Insurance rests on a balance between individual risk and collective protection, but contemporary pricing models based on massive datasets and opaque algorithms raise pressing fairness issues with respect to protected characteristics. While actuarial standards prioritize risk-based accuracy, stakeholders also demand ethical accountability, social solidarity, and protection against insidious discrimination. We place these three pillars within broader debates in business ethics and algorithmic fairness, linking them to traditions of distributive justice (Rawls (1971), Sen (1992)), information ethics (Floridi (2016), Nissenbaum (2009)), and risk-sharing theory (Arrow (1963)). We argue that these three pillars make explicit the ethical trade-offs actuaries and insurers face when deploying predictive models. No single principle can dominate without cost: actuarial fairness may increase socioeconomic disparities, solidarity may weaken market efficiency, and the causal approach relies on unverifiable assumptions that may reduce predictive power. This three-dimensional framework therefore turns fairness into an explicit governance objective, useful for regulation, corporate governance, and accountability in insurance. Beyond the actuarial domain, it also offers a generalizable framework for evaluating fairness in other risk-based algorithmic decisions, from credit scoring to healthcare pricing.

Perceived Fairness in Networks

by Arthur Charpentier, is available online at [doi:10.48550/arXiv.2510.12028](https://doi.org/10.48550/arXiv.2510.12028)

“Classical definitions of algorithmic fairness rely on aggregate statistics, such as demographic parity or equal opportunity. Yet in many social and economic settings, fairness is perceived locally, through peer networks and individual comparisons. We propose a theoretical model of perceived-fairness networks, in which the feeling of discrimination depends on the local topology of interactions. We show that, even when a decision rule satisfies standard fairness criteria, perceived discrimination may persist, or even intensify, due to homophily or assortative mixing. Analytical and simulation results thus show how network topology widens the gap between objective fairness and perceived fairness, with implications for algorithmic governance, finance, and collaborative insurance.


This paper revisits ideas that were presented during the **decentralized insurance and risk sharing** workshop,  at the **Insurance: Mathematics & Economics** conference, held in Chicago at the beginning of summer 2024.

The idea here is to use the topology induced by the graph. To keep things simple, when we have observations $\{y_1, \dots, y_n\}$ (or \mathbf{y}), the average over the whole group is

$$\bar{y} = \frac{1}{n} \sum_{j=1}^n y_j = \frac{\mathbf{1}^\top \mathbf{y}}{\mathbf{1}^\top \mathbf{1}}.$$


But if an individual is connected only to certain people through a graph \mathcal{G}_n , with adjacency matrix \mathbf{A} , the average observed over the d_i neighbors of individual i is

$$\bar{y}(i) = \frac{1}{d_i} \sum_{j \in N_i} y_j = \frac{1}{d_i} \sum_{i=1}^n A_{i,j} y_j = \frac{\mathbf{A}_i^\top \mathbf{y}}{\mathbf{A}_i^\top \mathbf{1}}.$$



The idea is to propose local versions of fairness metrics, no longer computed as averages over n people, but as local averages over each individual's connections.

In a graph with homophily (following the saying “birds of a feather flock together”), these quantities can be substantially different.



Sequential Transport for Causal Mediation Analysis

with Agathe Fernandes Machado, Iryna Voitsitska (Ірина Войціцька), Arthur Charpentier and Ewen Gallic, [doi:10.48550/arXiv.2603.15182](https://doi.org/10.48550/arXiv.2603.15182)

In causal inference, the counterfactual for one observation is the outcome that this observation would have had under a different condition from the one actually experienced. For example, if a person received a medical treatment and recovered, the counterfactual is what would have happened to that same person had they not received the treatment. Since only one scenario can be observed for each individual (treated or untreated), the counterfactual is always unobserved and must be estimated using models or comparisons with similar individuals. This estimation is essential to determine the individual causal effect of an intervention (\mathbb{W}).

The decomposition of the causal effect into direct and indirect effects, from Pearl (2001), where $x(a)$ denotes a “potential feature” (in x) for group a , and $\hat{\mu}_a$ is a predictive model for

individuals in group a , then for $a \in \{0, 1\}$

$$\tau := \delta(a) + \zeta(1 - a),$$

where


$$\begin{cases} \delta(a) := \widehat{\mu}_a(\mathbf{x}(1)) - \widehat{\mu}_a(\mathbf{x}(0)) \\ \zeta(a) := \widehat{\mu}_1(\mathbf{x}(a)) - \widehat{\mu}_0(\mathbf{x}(a)) \end{cases}$$

is reminiscent of the decomposition of Kitagawa (1955), Oaxaca (1973) and Blinder (1973)

$$\tau := \underbrace{\widehat{\beta}_a(\bar{x}_1 - \bar{x}_0)}_{\delta(a)} + \underbrace{(\widehat{\beta}_1 - \widehat{\beta}_0)\bar{x}_{1-a}}_{\zeta(1-a)}.$$

which we discussed in Newsletter #2. Instead of taking the average in the treated group, \bar{x}_1 , we propose an algorithm to generate a counterfactual value, $x(1)$.

“ We propose sequential transport (ST), a distributional framework for mediation analysis that combines optimal transport (OT) with a directed acyclic graph (DAG) of mediators. Instead of relying on cross-world counterfactual assumptions, ST constructs unit-level mediator counterfactuals by minimally transporting each mediator — marginally or conditionally — to its distribution under an alternative treatment, while preserving the causal dependencies encoded by the DAG. For numerical mediators, ST uses monotone (conditional) OT maps based on conditional CDF/quantile estimators; for categorical mediators, it extends naturally through simplex-based transport. We establish consistency of the estimated transport maps and of the induced unit-level decompositions into direct and indirect effects, *mutatis mutandis*, under standard regularity and support conditions. When treatment is randomized or ignorable (possibly conditional on covariates), these decompositions admit a causal interpretation. Gaussian examples show that ST recovers classical mediation formulas, while additional simulations confirm good performance in nonlinear and mixed settings. An application to the COMPAS dataset illustrates how ST produces deterministic mediator counterfactuals consistent with DAGs.

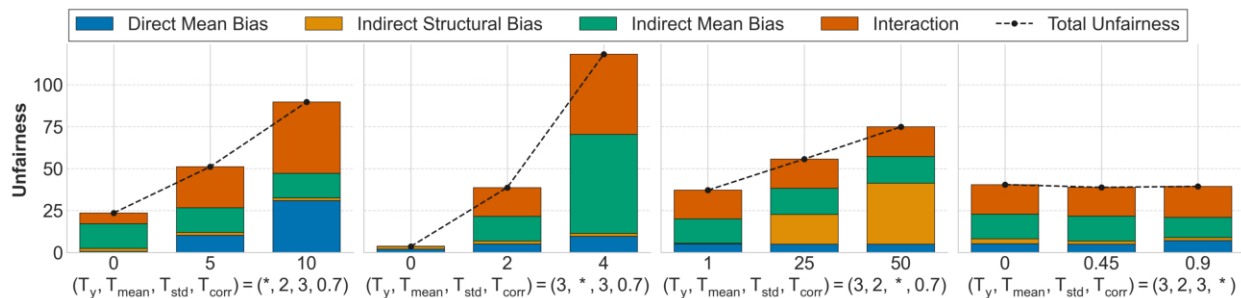


Criminal Records, the Right to be Forgotten, and Inclusive Insurance

by Arthur Charpentier and David Schraub, SSRN:6170208.

“ Criminal history information (criminal history, CHx) is increasingly used for insurance screening and underwriting, yet remains undertheorized in the inclusive-insurance literature, even though it may expand protection gaps (quote-stage blocking, opaque third-party data chains, limited recourse). This paper argues that CHx-based underwriting belongs both to risk governance and to information governance: criminal data that are legally and temporally structured can be transformed, through automation and outside vendors, into durable stigma — especially when copies and derived products persist without reflecting changes in legal status (sealed records, expungement, rehabilitation). The paper maps the entry points of CHx across the contract life cycle and highlights inclusion failures (provenance, quality, update delays, procedural safeguards). It then proposes a framework distinguishing insurance relevance from “digital punishment,” translated into operational mechanisms (time windows / decaying weights, minimization and purpose limitation, auditable vendor governance, and rapid notice–explanation–correction processes).

A French version is available on HAL, hal:05490251.



Decomposing Direct and Indirect Biases in Linear Models under Demographic Parity Constraint

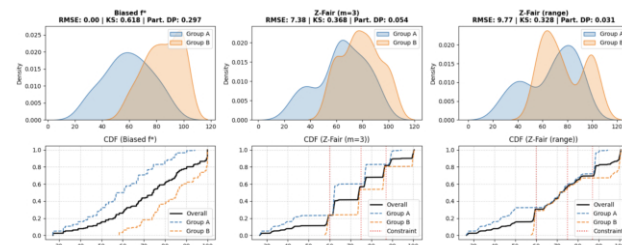
with Bertille Tierny, Arthur Charpentier and François Hu, available online at [doi:10.48550/arXiv.2511.11294](https://doi.org/10.48550/arXiv.2511.11294).

“Linear models are widely used in high-stakes decision-making because of their simplicity and interpretability. However, when fairness constraints such as demographic parity are introduced, their effects on model coefficients — and therefore on the distribution of predictive bias across features — remain opaque. Existing approaches for linear models often rely on strong and unrealistic assumptions, or neglect the explicit role of the sensitive attribute, which limits their practical usefulness for fairness assessment. We extend the work of Chzhen & Schreuder (2022) and Fukuchi & Sakuma (2023) by proposing a post-processing framework that can be applied to any linear model in order to decompose the resulting bias into direct components (sensitive attributes) and indirect components (correlated features). Our method analytically characterizes how demographic parity reshapes each model coefficient, including those of sensitive and non-sensitive features. This allows a transparent, feature-level interpretation of fairness interventions and reveals how bias can persist or shift through correlated variables. Our framework requires no re-training and provides actionable insights for model auditing and mitigation. Experiments on synthetic and real datasets show that our method captures fairness dynamics missed by previous work.

As discussed on page 12, this work was presented at the 40th Annual AAI Conference on Artificial Intelligence in Singapore.

Fair regression under localized demographic parity constraints

by Arthur Charpentier, Christophe Denis, Romuald Élie, Mohamed Hebiri and François Hu, available online at [hal:05567866](https://hal.archives-ouvertes.fr/hal-05567866).



“Demographic parity (DP) is a widely used group-fairness criterion that requires predictive distributions to be invariant across sensitive groups. While natural in classification, full distributional DP is often too restrictive in regression and can lead to a substantial loss of accuracy. We propose a relaxation of DP tailored to regression, by imposing parity only at a finite set of quantile levels and/or score thresholds. Concretely, we introduce a new fair predictor (ℓ, \mathcal{Z}) , which enforces group-wise CDF constraints of the form $F_{f|S=s}(z_m) = \ell_m$ for prescribed pairs (ℓ_m, z_m) . For this parametrization, we derive closed-form characterizations of the optimal discretized fair predictor via a Lagrangian dual formulation and quantify the cost of discretization, showing that the risk gap relative to the continuous optimum vanishes as the grid is refined.

We further develop a model-agnostic post-processing algorithm based on two samples (labeled for learning a base regressor and unlabeled for calibration), and establish finite-sample guarantees on constraint violation and excess penalized risk. In addition, we introduce two alternative frameworks in which we match marginal and group CDF values at selected score thresholds. In both cases, we provide closed-form solutions for the optimal discretized fair predictor. Experiments on synthetic and real datasets illustrate an interpretable trade-off between fairness and accuracy, allowing targeted corrections at decision-relevant quantiles or thresholds while preserving predictive performance.

Delving into Deep Smoothed Bootstrap: Application in Imbalanced Regression

by Samuel Stocksieker, Denys Pommeret and Arthur Charpentier, was presented at the [12th IEEE International Conference on Data Science and Advanced Analytics \(DSAA'25\)](#), held in Birmingham, England, in mid-October 2025.

“ Learning from an imbalanced distribution is a major and common challenge in predictive modeling. It usually leads to reduced performance of standard algorithms. Several approaches exist to address this issue, but many are tailored to classification problems, with limited relevance for regression. In this paper, we present a new method to improve learning on tabular data in the setting of imbalanced regression (IR), which remains an important problem. We propose using variational autoencoders (VAEs), which are known as a powerful tool for synthetic data generation. They offer an interesting approach for modeling and capturing latent representations of complex distributions. However, like all standard approaches, VAEs can be ineffective in the presence of data imbalance. We therefore develop an innovative data-generation approach combining VAEs with smoothed bootstrap.

KurtHGR: A Neural Maximal Correlation for Tabular Datasets

by Samuel Stocksieker, Denys Pommeret and Arthur Charpentier, online at [doi:10.1016/j.procs.2025.09.338](https://doi.org/10.1016/j.procs.2025.09.338),

“ The study of dependencies between variables is a fundamental pillar of machine learning, with implications for feature selection, fairness, dimensionality reduction, and multimodal learning. Among nonlinear correlation measures, Hirschfeld–Gebelein–Rényi (HGR) maximal correlation stands out for its universality and remarkable theoretical properties. Defined as the maximum achievable correlation between nonlinear transformations of two random variables, it provides an intrinsic quantification of statistical dependence, regardless of their marginal distributions. However, despite its theoretical potential, its practical adoption still faces several difficulties. In this article, we present a new approach called KurtHGR, dedicated to estimating the bivariate nonlinear correlation matrix of a set of variables. We show that this solution is effective for detecting nonlinear correlations, robust to noise, and computationally efficient, thanks to a neural architecture specifically designed for this task. We evaluate its performance through numerical illustrations and feature-selection experiments, where we show that KurtHGR empirically outperforms the current state of the art.

In a nutshell, maximal correlation is the greatest correlation (in the usual Pearson sense) that can be achieved by transforming variables in a nonlinear way. This search for the best nonlinear transformations makes it a fairly complex functional optimization problem, and neural networks are a promising approach.

Improving Procedural and Substantive Fairness in Conformal Prediction

by Pengqi Liu, Zijun Yu, Massoud Asgharian, Arthur Charpentier, and the TD Insurance and Layer 6 team, Jesse Cresswell and Mouloud Belbahri, online at [doi:10.48550/arXiv.2602.16794](https://doi.org/10.48550/arXiv.2602.16794)

“Conformal prediction (CP) provides distribution-free uncertainty quantification for machine-learning models, but its interaction with fairness in downstream decision-making remains little explored. Beyond CP as a stand-alone operation (procedural fairness), we analyze the holistic decision-making process in order to evaluate substantive fairness, that is, the fairness of downstream outcomes. On the theoretical side, we derive an upper bound that decomposes disparity in prediction-set size into interpretable components, thereby clarifying how label-conditional CP helps control the contribution of methods to unfairness. To facilitate scalable empirical analysis, we introduce an LLM-in-the-loop evaluator that approximates human evaluation of substantive fairness across several modalities. Our experiments show that label-conditional CP variants consistently deliver better substantive fairness. Finally, we show empirically that equalized set size, rather than coverage, is strongly correlated with improved substantive fairness, enabling practitioners to design fairer CP systems.

In conformal prediction, we do not return a single prediction, or a single “label” (in a classification problem), \hat{y} , but rather a set of plausible predictions $\mathcal{C}(x)$ (a “conformal set”) that is constructed to contain the true answer y with guaranteed probability.

The starting point of this paper was that procedural fairness refers to the “fairness” of the conformal sets themselves, independently of how they will be used: for example, one can look at whether coverage is comparable across groups (*equalized coverage*), or whether set size or the rate of singleton sets differs by group. By contrast, substantive fair-

ness concerns the real consequences of these sets in a decision pipeline: when a decision-maker (human/LLM) receives x and the set $\mathcal{C}(x)$, does this assistance improve performance (for example, *accuracy*) equally across all groups? The key idea is that a procedure may appear “fair” under procedural metrics (e.g. small coverage gaps) while still producing substantial *downstream* disparities, hence the need to evaluate substantive fairness as well.

The Python code is available online,

<https://github.com/layer6ai-labs/llm...>

The Impossible Right to Make Mistakes, the Impossible Right to Be Forgotten?

by Arthur Charpentier, published in [Risques](#), issue 144.

“We used to say that public information is like toothpaste; once it is out of the tube, it cannot be put back in. In a now fully digital world, every action leaves a trace, whether it is a deferred payment, an online search, or a geolocated trip. Aggregated together, these traces form risk profiles, creditworthiness scores, or lasting reputations. But a just society cannot be reduced to total memory: the right to make mistakes and the right to be forgotten are conditions for a second chance. Japanese wisdom teaches that one falls seven times and gets up eight (七転び八起き): human beings move forward through trial and error, provided that their mistakes are not frozen into a permanent stigma.

Incidentally, Arthur and David Schraub are participating in a working group in the United States on the use of criminal records by insurers (also an important issue in Canada, see [↪](#)). This article is the preliminary French version (and a shorter one) of the article presented on page 8.

What If We Finally Invested in Prevention?

by Arthur Charpentier, published in [Risques](#) (145)

“Prevention lies at the heart of any risk-management policy: it aims to reduce future costs by acting before a loss occurs. Health, climate, cybersecurity, or road safety: in every case, the benefits have been demonstrated, yet investment remains marginal. In insurance systems, this paradox is especially striking. Although insurers have much to gain from a lasting reduction in claims, their economic models continue to favor compensation over anticipation.”

Conference Presentations


Taking advantage of his time in Asia, Arthur Charpentier was able to give several talks in the region.

AAAI'26 (and visit) in Singapore

The 40th Annual AAI Conference on Artificial Intelligence was held in Singapore at the end of January (after Vancouver in 2024, [Newsletter #1](#) ↪, and Philadelphia in 2025, [Newsletter #3](#) ↪). AAI is rated an “A* conference” according to the [CORE ranking](#) ↪.



François first presented in the *main track*, in the session **ML: Ethics, Bias, and Fairness 1**. Then, Bertille Tierny ([Newsletter #4](#) ↪) was selected to present [Decomposing Direct and Indirect Biases in Linear Models under Demographic Parity Constraint](#) in the **AAAI Student** session, and Bertille was a finalist in the *3-Minute Thesis (3MT) Competition* for the **EGSAI Community Activity** ↪, “Bridging Gaps for Underrepresented Researchers in AI,” ↪



Arthur Charpentier stayed for two weeks at **ESSEC** (Asia-Pacific, Singapore campus), ↪, invited by **Pierre Alquier** (Professor of Statistics, Department of Information Systems, Data Analytics and Operations) ↪



Arthur gave a first talk, [Modeling and Understanding Indirect Discrimination in Algorithmic Fairness](#), at the ESSEC faculty seminar, and a second, [Balance and Calibration of Probabilistic Scores: From GLM to Machine Learning](#), at the **working group on risk**. He also gave a talk, [Fairness and Discrimination in Insurance](#), at the actuarial science seminar at **Nanyang Technological University** (NTU), invited by Ian Weng (Sophia) Chan.


Seminars and Conferences

- Agathe Fernandes Machado gave a talk at **Polytechnique Montréal** on [Optimal Transport for Counterfactual Fairness](#), invited by Sébastien Gambs.
- Arthur Charpentier gave a talk, [Detecting Hidden Biases in Insurance AI with Counterfactual Methods](#), at the **III Congreso Universitario Internacional sobre Seguros y Reaseguros**, ↪, organized in Lima, Peru (in hybrid format).





- Olivier Côté was in Switzerland, at the **Eidgenössische Technische Hochschule Zürich** (ETHZ), this winter for the first edition of the **ASTIN Bulletin Conference**, where he presented [Exposing Indirect Discrimination](#), ↪. At the same conference, Sang Huynh presented joint work with Le Ngan Dao, Van Yen Vuong, and Arthur Charpentier, [A Transport-based Framework](#)


- for Fair Insurance Premium Across Multiple Coverages.
- François Hu was invited to the **Actuarial Technology Summit 2025** to present [Bias Testing and the U.S. Regulatory Framework](#), based on work with Bertille Tierny and Arthur, presented on page 9.
 - Agathe gave a talk, [Assessing Counterfactual Fairness via \(Marginally\) Optimal Transport](#), at an **Algorithmic Fairness and Causality** event in Paris, invited by Marianne Clausel.
 - Olivier presented [Profiling Actuarial Discrimination via Causal Decomposition](#) at the **Actuarial and Financial Mathematics Conference**, held in Brussels this winter 🌐. A year earlier ([Newsletter #3](#) ➡), Agathe and Marie-Pier had also presented our work on fairness there.
 - Marie-Pier Côté was invited to present at the 8th edition of the **Workshop on Insurance Mathematics (WIM)**, held at the University of Toronto, where she presented [Unfounded Prediction Decomposition](#), based on work with Olivier, Marouane, and Arthur.
 - Olivier presented [Operationalizing Fairness in Actuarial Pricing: From Principle to Practice](#) at the **CAS Annual Meeting** of the **Casualty Actuarial Society**, 🌐 in Austin, Texas, in a session with Mallika Bender.
 - Arthur presented [Counterfactual and Transport-Based Methods for Understanding Indirect Discrimination in Algorithmic Systems](#) at the **JFLI Seminar** organized at the 国立情報学研究所 (National Institute of Informatics), in Tokyo.
 - Olivier was invited to spend a few days at **Stockholms universitet**, in Sweden, by Filip Lindskog and Mathias Lindholm, where he gave a seminar on [Profiling Actuarial Discrimination via Causal Decomposition](#).
 - Marouane Il-Idrissi presented [Interpretation of Black-Box Predictive Models](#) at the **StatQAM statistics seminar**, in Montreal.
 - Arthur gave a talk, [Actuarial Pricing Discrimination and Fairness](#), to more than one hundred actuaries in the AXA group during a **pricing seminar**, and the following week [Balance and Calibration in Insurance Pricing](#) during a **technical seminar**.
 - Marouane Il-Idrissi presented [Interpretability and Explainability of Black-Box Methods](#) at the **Dalhousie Statistics Seminar**, 🌐, in Halifax, Nova Scotia.
 - Agathe presented [Assessing Counterfactual Fairness via \(Marginally\) Optimal Transport](#) at the **AMSE econometrics seminar** (Aix-Marseille School of Economics), 🌐.
 - Olivier was invited to present at the conference **Bias in Assessing Financial Risk: Insights from the 2025 Research Award** of the **American Academy of Actuaries**. He gave a presentation on [Exposing Indirect Discrimination](#), and took part in the panel discussion that followed, with more than 400 participants, 🌐.






- Agathe presented [Assessing Counterfactual Fairness via \(Marginally\) Optimal Transport Theory](#) at the SIAS seminar of the computer science department of **University of Amsterdam**, invited by (Fernando) Pascoal Dos Santos 🌐, who had been her mentor for the *Doctoral Consortium* at IJCAI last summer.
- Agathe was also invited to **Erasmus University Rotterdam** by Aurelie Lemmens, where she presented [Assessing Counterfactual Fairness via \(Marginally\) Optimal Transport Theory](#).
- Marouane and Lucas Offroy jointly presented [Cooperative Games for the Interpretation of Machine Learning Models](#) at the **PSPP Seminar**, 🌐, at EDF, in France.
- Agathe was invited to take part in the workshop **AI in Risk Assessment and Mitigation**, held at ICMS, Bayes Centre, Edinburgh, 🌐, for an event organized by Marie Michaelides, George Streftaris, and George Tzougas (**Heriot-Watt University**). She gave a talk on [Measuring Fairness through Calibration in Insurance Scoring Algorithms](#).

- Marouane was invited to present at the **CIRRELT Reading Club** (Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport) — **Optimization and Game Theory** at GERAD (Groupe d'études et de recherche en analyse des décisions). He presented [Beyond Shapley Values: Cooperative Games for the Interpretation of Machine Learning Models](#).
- Arthur gave a short talk on [From Climate to Insured Losses \(Sources of Uncertainty\)](#) in a working group at the **Institut Louis Bachelier**.
- Marouane organized the [Explainable AI](#) session at the **MobiliT.AI'25** conference, , co-organized by IVADO, IRT Saint Exupéry, IID, CRIAQ, and ANITI.
- Olivier presented [Operationalizing Fairness in Actuarial Pricing: From Principle to Practice](#) at the **2025 CAS Canada Connection**, , in Toronto, Canada.
- Agathe gave a talk, [Shapley Values for Conformal Prediction Intervals](#), at the **UBO mathematics seminar** (Université de Bretagne Occidentale), in Brest.
- Agathe gave a talk on [Revealing Sources of Uncertainty in Machine Learning Models: Contribution of Explanatory Variables to Conformal Prediction Intervals](#) at the **Quantact Seminar**, of the actuarial science team, at **Université Laval**, in Quebec City. She also took the opportunity to work with Olivier and Marie-Pier Côté. We will return to that soon...



- Marouane was invited to take part in the **SIAM Conference on Uncertainty Quantification**, held in March in Minneapolis, Minnesota, U.S. . Given recent events, it was decided not to participate, for obvious safety reasons.
- Marouane presented [Cooperative Games for the Interpretation of Machine Learning](#)

[Models](#) at the **AMSE econometrics seminar** (Aix-Marseille School of Economics).

- In early December, Benoît Régent-Kloeckner was invited to present [Climate Updating of Flood Risk in Metropolitan France](#) at an event of the **Institut des Actuaire**s, in France.
- Marouane presented [Cooperative Games for the Interpretation of Machine Learning Models](#) at the Statistics Seminar of Université Laval, in Quebec City.
- Olivier presented [A Scalable Toolbox for Exposing Indirect Discrimination in Insurance Rates](#) to the TD Insurance teams in Montreal.
- Marouane gave a presentation on [Generalizing SHAP Values](#) during the **IntactLab Methodology Meetings** in November, in Montreal.
- Agathe gave a talk on [Assessing Counterfactual Fairness via Optimal Transport Theory](#) for the R&D teams at **JPMorganChase**.
- Marouane presented [Cooperative Games for the Interpretation of Machine Learning Models](#) at the **IRIT UT Capitole Seminar**, , in Toulouse, France.
- Stéphane Mussard was in Bologna, Italy, this fall to present [KNN and K-Means in Gini Prametric Spaces](#) at **ECAI 2025 – 28th European Conference on Artificial Intelligence**, , rated an “A” conference according to the [CORE ranking](#) .



- Marouane presented [Cooperative Games for the Interpretation of Machine Learning Models](#) at the **Statistics Seminar of the University of Sherbrooke**, in Canada.
- Denys Pommeret presented [Delving into Deep Smoothed Bootstrap: Application in Imbalanced Regression](#) at the **IEEE DSAA (12th IEEE International Conference on**

Data Science and Advanced Analytics) 🌐, held in Birmingham, UK.

- Arthur was invited to the **10th Annual Global Congress of Knowledge Economy** (GCKE), held in November in Osaka, 大阪市 🌐. He was in charge of session GCKE04, Economic Governance & Sustainable Development. He gave a talk on *Fairness and Insurance: Disentangling Illegitimate and Indirect Discriminations*.
- Raphaël Dalbarade presented *Rate Segmentation in Natural Catastrophes: Issues and a Stackelberg Model* at the **P&C Day of the Institut des Actuares**, 🌐, in France, with Laurence Barry. The day focused on “*segmentation and pooling: theory, practice and future perspectives.*”



The following day, François Hu presented *Data, Models, and Fairness: How Should We Address Bias in Pricing Models?*, based on work presented in earlier newsletters.

- This fall, Ewen Gallic presented *Assessment of Counterfactual Fairness Using Marginally Optimal Transport* at the **Economics Seminar of ENSAI**, the National School for Statistics, in Rennes, France.
- Agathe presented *Causal Mediation Analysis via Sequential Transport to Assess Counterfactual Fairness* at the conference **Mathematical Foundations of AI – When Uncertainty Quantification Meets Causal Modelling and Inference**, 🌐, in Paris, organized by the Research Theme in Uncertainty Quantification (CNRS), the CAUSALITAI project, and the AI clusters MIAI, SCAI, and DATAIA.
- Agathe took part in the **MLISTRAL2** conference (*Machine Learning in Insurance Sector Targeted to Risk Analysis and Losses*), 🌐, at **CIRM** (Centre international de rencontres mathématiques) in Marseille, co-organized by Samuel Stocksieker, Katrien Antonio, Xavier Milhau, and Mathias Valla.
- Arthur presented *A Three-Pillar Governance*

Framework for Fair Insurance Pricing during a seminar of the **PARI Chair**, online.

Ten years after the Paris Agreement, the conference *Green Finance Research Advances 2025* was organized by ILB and the PARC group, Paris Agreement Research Commons (of which Arthur is a member).



🕒 Coming Up

- This spring, Bei Jiang and Linglong Kong (University of Alberta) are organizing a 5-day workshop, “**Statistics in Trustworthy AI**,” during an IVADO thematic session, *Thematic Program on Statistical Foundations of AI*. Marouane will take part, leading a workshop on explainability. He has also been invited to speak at the annual conference of the Statistical Society of Canada, **SSC**.
- Agathe has been invited to spend a little more than 2 months at **UNSW**, Sydney, Australia, by **Fei Wang** (Newsletter #2 🔄). On the way there, she will stop in Kyoto, 京大, in Japan.
- Next month, Arthur will be at **Tsinghua University** 清华大学, in Beijing, 海淀区, China, invited by Runhuan Feng 冯润桓, for the joint seminar of **Renmin University** 中国人民大学 and **Peking University** 北京大学, then at the **Chinese University of Hong Kong**, 香港中文大學, invited by Phillip Yam 任尚智.

Newsletters #1 🔄 #2 🔄 #3 🔄 #4 🔄

SCOR
FOUNDATION FOR SCIENCE

Newsletter #5/6

Project funded by SCOR Foundation for Science