LIST OF ABSTRACTS . 2024

2024 / 01

Affine Heston model style with self-exciting jumps and long memory

Charles Guy Leunga Njike, Donatien Hainaut

Classic diffusion processes fail to explain asset return volatility. Many empirical findings on asset return time series, such as heavy tails, skewness and volatility clustering, suggest decomposing the volatility of an asset's return into two components, one caused by a Brownian motion and another by a jump process. We analyze the sensitivity of European call options to memory and self-excitation parameters, underlying price, volatility and jump risks. We expand Heston's stochastic volatility model by adding to the instantaneous asset prices, a jump component driven by a Hawkes process with a kernel function or memory kernel that is a Fourier transform of a probability measure. This kernel function defines the memory of the asset price process. For instance, if it is fast decreasing, the contagion effect between asset price jumps is limited in time. Otherwise, the processes remember the history of asset price jumps for a long period. To investigate the impact of different rates of decay or types of memory, we consider four probability measures: Laplace, Gaussian, Logistic and Cauchy. Unlike Hawkes processes with exponential kernels, the Markov property is lost but stationarity is preserved; this ensures that the unconditional expected arrival rate of the jump does not explode. In the absence of the Markov property, we use the Fourier transform representation to derive a closed form expression of a European call option price based on characteristic functions. A numerical illustration shows that our extension of the Heston model achieves a better fit of the Euro Stoxx 50 option data than the standard version.

2024 / 02

Option pricing in the Heston model with Physics inspired neural networks

Donatien Hainaut, Alex Casas

In absence of a closed form expression such as in the Heston model, the option pricing is computationally intensive when calibrating a model to market quotes. this article proposes an alternative to standard pricing methods based on physics-inspired neural networks (PINNs). A PINN integrates principles from physics into its learning process to enhance its efficiency in solving complex problems. In this article, the driving principle is the Feynman-Kac (FK) equation, which is a partial differential equation (PDE) governing the derivative price in the Heston model. We focus on the valuation of European options and show that PINNs constitute an efficient alternative for pricing options with various specifications and parameters without the need for retraining.

2024 / 03

Sliced-Wasserstein Estimation with Spherical Harmonics as Control Variates

Rémi Leluc, Aymeric Dieuleveut, François Portier, Johan Segers, Aigerim Zhuman

The Sliced-Wasserstein (SW) distance between probability measures is defined as the average of the Wasserstein distances resulting for the associated one-dimensional projections. As a consequence, the SW distance can be written as an integral with respect to the uniform measure on the sphere and the Monte Carlo framework can be employed for calculating the SW distance. Spherical harmonics are polynomials on the sphere that form an orthonormal basis of the set of square-integrable functions on the sphere. Putting these two facts together, a new Monte Carlo method, hereby referred to as Spherical Harmonics Control Variates (SHCV), is proposed for approximating the SW distance using spherical harmonics as control variates. The resulting approach is shown to have good theoretical properties, e.g., a no-error property for Gaussian measures under a certain form of linear dependency between the variables. Moreover, an improved rate of convergence, compared to Monte Carlo, is established for general measures. The convergence analysis relies on the Lipschitz property associated to the SW integrand. Several numerical experiments demonstrate the superior performance of SHCV against state-of-the-art methods for SW distance computation.

2024 / 04

Nonlinear wavelet threshold estimation of time-varying covariance matrices in a log-Euclidean manifold

Gabriel Bailly, Rainer von Sachs

We tackle the problem of estimating time-varying covariance matrices (TVCM; i.e. covariance matrices with entries being time-dependent curves) whose elements show inhomogeneous smoothness over time (e.g. pronounced local peaks). To address this challenge, wavelet denoising estimators are particularly appropriate. Specifically, we model TVCM using a signal-noise model within the Riemannian manifold of symmetric positive definite matrices (endowed with the log-Euclidean metric) and use the intrinsic wavelet transform, designed for curves in Riemannian manifolds. Within this non-Euclidean framework, the proposed estimators preserve positive definiteness. Although linear wavelet estimators for smooth TVCM achieve good results in various scenarios, they are less suitable if the underlying curve features singularities. Consequently, our estimator is designed around a nonlinear thresholding scheme, tailored to the characteristics of the noise in covariance matrix regression models. The effectiveness of this novel nonlinear scheme is assessed by deriving mean-squared error consistency and by numerical simulations, and its practical application is demonstrated on TVCM of electroencephalography (EEG) data showing abrupt transients over time.

2024 / 05

A penalised bootstrap estimation procedure for the explained Gini coefficient

Alexandre Jacquemain, Cédric Heuchenne, Eugen Pircalabelu

The Lorenz regression estimates the explained Gini coefficient, a quantity with a natural application in the measurement of inequality of opportunity. Assuming a single-index model, it corresponds to the Gini coefficient of the conditional expectation of a response given some covariates and it can be estimated without having to estimate the link function. However, it is prone to overestimation when many covariates are included. In this paper, we propose a penalised bootstrap procedure which selects the relevant covariates and produces valid inference for the explained Gini coefficient. The obtained estimator achieves the Oracle property. Numerically, it is computed by the SCAD-FABS algorithm, an adaptation of the FABS algorithm to the SCAD penalty. The performance of the procedure is ensured by theoretical guarantees and assessed via Monte-Carlo simulations. Finally, a real data example is presented.

2024 / 06

Conditional expectations given the sum of independent random variables with regularly varying densities

Michel Denuit, Patricia Ortega-Jimenez, Christian Y. Robert

Stochastic monotonicity of two independent random variables X and Y given the value of their sum S = X + Y has been linked to log-concave densities since Efron (1965). However, the log-concavity assumption is not realistic in some applications because it excludes heavy- tailed distributions. This paper considers random variables with regularly varying densities to illustrate how heavy tails can lead to a non-monotonic behavior for the conditional expectation mX(s) = E[X|S = s], which turns out to be problematic in risk sharing or signal processing (including industry loss warranties or parametric insurance, for instance). This paper first aims to identify situations where a non-monotonic behavior appears according to the tail-heaviness of X and Y . Secondly the paper aims to study the asymptotic behavior of mX (s) as the value s of the sum gets large. The analysis is then extended to zero-augmented probability distributions, commonly encountered in applications to insurance and to sums of more than two random variables. Consequences for signal processing and risk sharing are discussed. Many numerical examples illustrate the results.

2024 / 07

The effect of stock splits on liquidity in a dynamic model

Christian Hafner, Oliver Linton, Linqi Wang

We develop a dynamic framework to detect the occurrence of permanent and transitory breaks in the illiquidity process. We propose various tests that can be applied separately to individual events and can be aggregated across different events over time for a given firm or across different firms. In an empirical study, we use this methodology to study the impact of stock splits on the illiquidity dynamics of the Dow Jones index constituents and the effects of reverse splits using stocks from the S&P 500, S&P 400 and S&P 600 indices. Our empirical results show that stock splits have a positive and significant effect on the permanent component of the illiquidity process while a majority of the stocks engaging in reverse splits experience an improvement in liquidity conditions.

2024 / 08

Latent Dirichlet Allocation for structured insurance data

Charlotte Jamotton, Donatien Hainaut

This article explores the application of Latent Dirichlet Allocation (LDA) to structured tabular insurance data. LDA is a probabilistic topic modelling approach initially developed in Natural Language Processing (NLP) to uncover the underlying structure of (unstructured) textual data. It was designed to represent textual documents as mixture of latent (hidden) topics, and topics as mixtures of words. This study introduces the LDA's document-topic distribution as a soft clustering tool for unsupervised learningtasks in the actuarial field. By defining each topic as a risk profile, and by treating insurance policies as documents and the modalities of categorical covariates as words, we show how LDA can be extended beyond textual data and can offer a framework to uncover underlying structures within insurance portfolios. Our experimental results and analysis highlight how the modelling of policies based on topic cluster membership, and the identification of dominant modalities within each risk profile, can give insights into the prominent risk factors contributing to higher or lower claim frequencies.

2024 / 09

Conical FDH Estimators of Directional Distances and Luenberger Productivity Indices for General Technologies

Cinzia Daraio, Simone Di Leo, Léopold Simar

In productivity and efficiency analysis, directional distances are very popular, due to their flexibility for choosing the direction to evaluate the distance of Decision Making Units (DMUs) to the efficient frontier of the production set. The theory and the statistical properties of these measures are today well known in various situations. But so far, the way to measure directional distances to the cone spanned by the attainable set has not been analyzed. In this paper we fill this gap and describe how to define and estimate directional distances to this cone, for general technologies, i.e. without imposing convexity. Their statistical properties are also developed. This allows us to measure distances to non-convex attainable set under Constant Returns to Scale (CRS) but also to measure and estimate Luenberger productivity indices and their decompositions for general technologies. The way to make inference on these indices is also described in details. We propose illustrations with some simulated data, as well as, a practical example of inference on Luenberger productivity indices and their decompositions with a well-known real data set.

2024 / 10

Central Limit Theorems for Directional Distance Functions with and without Undesirable Outputs

Léopold Simar, Valentin Zelenyuk, Shirong Zhao

We develop new central limit theorems (CLTs) for the aggregate directional distance functions (DDFs), which embed the CLTs for the aggregate efficiency and simple mean DDFs as special cases. Moreover, we develop new CLTs for the aggregate DDFs in the presence of the weak disposability of undesirable outputs. Our Monte-Carlo simulations confirm the good performance of statistical inference based on the new CLTs we have derived and illustrate how wrong the inference based on the standard CLTs can be. To our knowledge, this is the first study that provides both the asymptotic theory and the simulation evidence for the non-parametric frontier approaches when some outputs are undesirable. Finally, we provide an empirical illustration using a data set from large US banks as well as supply the computational code for alternative applications.

2024 / 11

Intergenerational risk sharing in pay-as-you-go pension schemes

Hélène Morsomme, Jennifer Alonso-Garcia, Pierre Devolder

Population ageing undermines traditional social security pension systems that combine pay-as-you-go (PAYG) and defined benefits (DB). Indeed, demographic risk, if guaranteed benefits remain unaltered, will be borne entirely by workers through increases in the contribution rate. To avoid a substantial increase of the contributions and in order to maintain simultaneously the financial sustainability and the social adequacy of the public pension system, risk sharing and automatic balancing mechanisms need to be put in place. We present a two-step convex family of risk-sharing mechanisms. The first shares the risk between contributors and retirees through adjustments in the contribution rate, used to calculate the global covered wage bill, and the benefit ratio that represents the relationship between average pensions and wages. The second step studies how the retirees' risk should be shared between the different retirees' generations through adjustments in the replacement rate and a sustainability factor that affects pension indexation during retirement. We perform a detailed study of the effect of social planner's targets and solidarity weight between various generations in a deterministic and stochastic environment.

2024 / 12

A Fast Method for Implementing Hypothesis Tests with Multiple Sample Splits in Nonparametric Models of Production

Léopold Simar, Paul Wilson

Kneip, Simar and Wilson (Journal of Business and Economic Statistics, 2016) and Daraio, Simar and Wilson (The Econometrics Journal, 2018) provide non-parametric tests of (i) convexity versus non-convexity of the production set, (ii) constant ver- sus non-constant returns-to-scale of the frontier, and (iii) separability versus non- separability of the frontier with respect to environmental variables. Among other uses, these tests are essential for deciding which non-parametric efficiency estimator should be used to estimate technical efficiency. Each test requires randomly splitting the sample. Although theory establishes that the tests are valid for any random split, results can vary with different splits. This paper provides a computationally efficient method to aggregate test outcomes across multiple sample-splits using ideas from the statistical literature on controlling false discovery rates in multiple testing situations. We provide tests using multiple sample-splits (to remove the ambiguity resulting from a single sample-split) and extensive Monte Carlo evidence on the size and power of our tests. The computational time required by the new tests is about 0.001 times the computational time required by the bootstrap method proposed by Simar and Wilson (Journal of Productivity Analysis, 2020).

2024 / 13

Efficient hedging of life insurance portfolio for loss-averse insurers

Edouard Motte, Donatien Hainaut

This paper investigates the hedging of equity-linked life insurance portfolio for loss-averse insurers. We consider a general arbitrage-free financial market and an actuarial market composed of n−independent policyholders. As the combined market is incomplete, perfect hedging of any actuarial-financial payoff is not possible. Instead, we study the efficient hedging of n−size equity-linked life insurance portfolio for insurers who are only concerned with their losses. To this end, we consider stochastic control problems (under the real-world measure) in order to determine the optimal hedging strategies that either maximize the probability of successful hedge (called quantile hedging) or minimize the expectation for a class of shortfall loss functions (called shortfall hedging). We show that the optimal strategies depend both on actuarial and financial risks. Moreover, these strategies adapt not only to the size of the insurance portfolio but also to the risk-aversion of the insurer. Under the additional assumption of complete financial market, we derive the explicit forms of the optimal hedging strategies. The numerical results show that, for loss-averse insurers, the optimal strategies outperform the optimal mean-variance hedging strategy, demonstrating the relevance of adopting the optimal strategy according to the insurers' risk aversion and portfolio size.

2024 / 14

Time-varying degree-corrected stochastic block models

Mengxue Li, Rainer von Sachs, Eugen Pircalabelu

Recent interest has emerged in community detection for dynamic networks which are observed along a trajectory of points in time. In this paper, we present a time-varying degree-corrected stochastic block model to fit a dynamic network which allows evolving heterogeneity in the degrees of nodes within a community over time. Considering the influence of the varying time window on the aggregation of network information from different time points, in the parameter estimation, we propose a smoothing-based method to recover time-varying degree parameters and communities. We also provide rates of consistency of our smoothed estimators for degree parameters and communities using a time-localised profile- likelihood approach. Extensive simulation studies and applications to two different real data sets complete our work.

2024 / 15

Participating life insurances in an equity-Libor Market Model

Donatien Hainaut, Laurent Devineau

This article introduces an Equity-Libor Market Model (LMM) that integrates the investment strategy into the valuation process of participating life insurance. Within this framework, we provide a semi-analytical formula for approximating the fair value of liabilities. We then investigate the impact of the investment policy on the net asset value and the solvency capital requirement. To carry out this analysis, we propose a Monte Carlo method for generating sample paths under both Libor and real measures, alongside an estimation procedure under the real measure. The numerical illustration focuses on the asset-liability management of an endowment and a life annuity.

2024 / 16

Deep learning for high-dimensional continuous-time stochastic optimal control without explicit solution

Jean-Loup Dupret, Donatien Hainaut

This paper introduces the GPI-PINN algorithm, a novel numerical scheme for solving continuous-time stochastic optimal control problems in high dimensions when the optimal control does not admit an explicit solution. Combining Physics-Informed Neural Networks with an Actor-Critic structure built upon the Generalized Policy Iteration technique, this successive deep learning algorithm employs two separate neural networks to approximate both the value function and the multidimensional optimal control. This way, the GPI-PINN algorithm manages to achieve a global approximation of the optimal solution across all time and space, which can be evaluated online rapidly. The optimality and convergence of the scheme are demonstrated theoretically and its accuracy and efficacy are shown empirically based on two numerical examples. In particular, we generalize the standard Almgren-Chriss model arising from optimal liquidation in finance by allowing for a price impact model with fully nonlinear temporary and permanent impact functions and by considering a multidimensional setting with numerous co-integrated assets.

2024 / 17

Risk times in mission-oriented systems

Antonio Arriaza, Jorge Navarro, Patricia Ortega-Jiménez

This article assesses risk times in mission-oriented systems with high safety standards. We examine critical times under two safety policies. The first requires that the system's reliability function, known the first failure of the components, must exceed a predetermined reliability level throughout the mission. The second demands that the conditional distribution of the system's lifetime, known the first failure of the components, must be more reliable, in terms of the usual stochastic order, than the original system's lifetime. Our study analyzes the critical times at which both policies remain viable. Our methodology, applicable to multiple failure scenarios, identifies sufficient conditions for the existence of these times. We offer explicit solutions for parallel systems with IID components and a general method for dependent and identically distributed components. The study includes practical examples and introduces a nonparametric estimator for the critical times.

2024 / 18

Tail calibration of probabilistic forecasts

Sam Allen, Jonathan Koh, Johan Segers, Johanna Ziegel

Probabilistic forecasts comprehensively describe the uncertainty in the unknown future outcome, making them essential for decision making and risk management. While several methods have been introduced to evaluate probabilistic forecasts, existing evaluation techniques are ill-suited to the evaluation of tail properties of such forecasts. However, these tail properties are often of particular interest to forecast users due to the severe impacts caused by extreme outcomes. In this work, we introduce a general notion of tail calibration for probabilistic forecasts, which allows forecasters to assess the reliability of their predictions for extreme outcomes. We study the relationships between tail calibration and standard notions of forecast calibration, and discuss connections to peaks-over-threshold models in extreme value theory. Diagnostic tools are introduced and applied in a case study on European precipitation forecasts.

2024 / 19

No-sabotage under conditional mean risk sharing of dependent-by-mixture insurance losses

Michel Denuit, Patricia Ortega-Jimenez, Christian Y. Robert

Conditional mean risk sharing defines an allocation rule to distribute total losses among participants in an insurance pool. Under this risk-sharing scheme, the no-sabotage condition holds when conditional expectations of individual losses given their sum are comonotonic. This property has been widely studied considering independent risks, often assuming that they possess log-concave densities. This paper considers the no-sabotage condition for dependent-by-mixture risks which do not necessarily obey log-concave distributions. Sufficient conditions derived from three different approaches are proposed in order to fulfill the no-sabotage requirement. Several examples are given to illustrate the applicability of the results.

2024 / 20

A panel analysis of microfinance efficiency measures: Evidence on the effects of unobserved managerial ability

François Seck Fall, Hubert Tchakoute Tchuigoua, Anne Vanhems, Léopold Simar

This study analyzes the impact of unobserved heterogeneity on microfinance efficiency. Using panel data for 168 microfinance institutions (MFIs) over the period 20102015, we examine the persistence of the effect of unobserved heterogeneity on microfinance efficiency. Using recent nonparametric and robust techniques, we identify a latent heterogeneity factor related to the ability of MFI managers to promote efficiency, independent of MFI size, and analyze its impact on MFI inefficiency measures over time. We then assess the robustness of our results to several factors: the MFI status of the MFI (for-profit or nonprofit), the definition of the efficiency measure (social and financial) and an observed degree of heterogeneity captured by the percentage of women on the board. Finally, we analyze the relationship between our unobserved heterogeneity factor and institutional and socio-economic indicators.

2024 / 21

European option pricing with model constrained Gaussian process regressions

Donatien Hainaut, Frédéric Vrins

We propose a method for pricing European options based on Gaussian processes. We convert the problem of solving the Feynman-Kac (FK) partial differential equation (PDE) into a model-constrained regression. We form two training sets by sampling state variables from the PDEs inner domain and terminal boundary. The regression function is then estimated to fit the option payoffs on the boundary sample while satisfying the FK PDE on the inner sample. We adopt a Bayesian framework in which payoffs and the value of the FK PDE in the boundary and inner samples are noised. Assuming the regression function is a Gaussian process, we find a closed- form approximation for the option prices. We demonstrate the performance of the procedure on call options in the Heston model and basket call options in a Black- Scholes market.

2024 / 22

Asymmetric Models for Realized Covariances

Luc Bauwens, Emilija Dzuverovic, Christian Hafner

We introduce asymmetric effects in the BEKK-type conditional autoregressive Wishart model for realized covariance matrices. The asymmetry terms are specified either by interacting the lagged realized covariances with the signs of the lagged daily returns or by using the decomposition of the lagged realized covariance matrix into positive, negative, and mixed semi-covariances, thus relying on the lagged intra-daily returns and their signs. We provide a detailed comparison of models with different complexity, for example with respect to restrictions on the parameter matrices. In an extensive empirical study, our results suggest that the asymmetric models outperform the symmetric one in terms of statistical and economic criteria. The asymmetric models using the signs of the daily returns tend to have a better in-sample fit and out-of-sample predictive ability than the models using the signed intra-daily returns.

2024 / 23

American option pricing with model constrained Gaussian process regressions

Donatien Hainaut

This article introduces a novel method based on Gaussian process regression for pricing American options. The variational partial differential equation (PDE) governing option prices is converted into a non-linear penalized Feynman-Kac equation (PFK). We propose an iterative algorithm to manage the non-linearity of the PFK operator. We sample state variables in the PDE's inner domain and on the terminal boundary. At each step, we fit a constrained regression function approximating the option price. This function matches the option payoffs on the boundary sample while satisfying the PFK PDE on the inner sample. The non-linear term in this PDE is frozen and valued with the price estimate from the previous iteration. We adopt a Bayesian framework in which payoffs and the value of the FK PDE in the boundary and inner samples are noised. Assuming the regression function is a Gaussian process, we find a closed-form approximation of option prices. In the numerical illustration, we evaluate American put options in the Heston model and in the two-factor Hull-White model.

2024 / 24

Bayesian mortality modelling with pandemics: a vanishing jump approach

Julius Goes, Karim Barigou, Anne Leucht

This paper extends the Lee-Carter model for single- and multi-populations to account for pandemic jump effects of vanishing kind, allowing for a more comprehensive and accurate representation of mortality rates during a pandemic, characterised by a high impact at the beginning and gradually vanishing effects over subsequent periods. While the Lee-Carter model is effective in capturing mortality trends, it may not be able to account for large, unexpected jumps in mortality rates caused by pandemics or wars. Existing models allow either for transient jumps with an effect of one period only or persistent jumps. However, there is no literature on estimating mortality time series with jumps having an effect over a small number of periods as typically observed in pandemics. The Bayesian approach allows to quantify the uncertainty around the parameter estimates. Empirical data from the COVID-19 pandemic shows the superiority of the proposed approach, compared to models with a transitory shock effect.

2024 / 25

Comparison of predictors' performance in insurance pricing: testing for Bregman dominance based on Murphy diagrams

Michel Denuit, Julien Trufin

Ehm et al. (2016) defined forecast dominance, or Bregman dominance as dominance for every Bregman loss function. This letter explores Bregman dominance to compare competing candidate pure premiums. An effective testing procedure for Bregman dominance is proposed based on Murphy diagrams and its performance is evaluated through a simulation study. An application to a Swiss motor insurance data set demonstrates the potential of the proposed procedure.