Probability, Finance and Signal

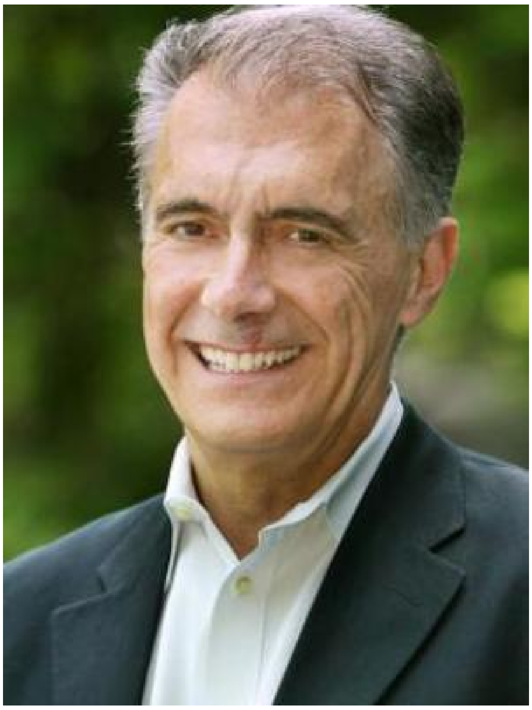

Conference in Honor of René Carmona

CIRM, Luminy, 19–23 May, 2025

BOOKLET OF ABSTRACTS

Saddlepoint Approximations for Hawkes Jump-Diffusion Processes with an Application to Risk Management

Yacine Aït-Sahalia (1), Roger J. A. Laeven (2),

(1) Department of Economics, Princeton University and NBER.
(2) Department of Quantitative Economics, University of Amsterdam.

We propose a statistical model based on Hawkes processes in which large financial losses can arise in close succession serially as well as cross-sectionally. We derive in closed form saddlepoint approximations to the tails of profit and loss distributions, both marginal and joint, and use them to construct explicit risk measure formulae that account for the fact that a given financial institution's losses make it more likely that that institution will experience further losses, and that other financial institutions will experience losses as well. These closed-form risk measures can be used for comparative statics, parameter calibration, and setting capital requirements and potential systemic risk charges.

Applications to Energy Savings : A Rank-Based Reward between a Principal and a Field of Agents

Clémence Alasseur (1),
Erhan Bayraktar (2),
Roxana Dumitrescu (3),
Quentin Jacquet (1),

(1) EDF Lab Saclay, Palaiseau / FIME research center, France.
(2) Department of Mathematics, University of Michigan, USA.
(3) ENSAE-CREST, Institut Polytechnique de Paris / FIME research center, France.

In this paper, we consider the problem of a Principal aiming at designing a reward function for a population of heterogeneous agents. We construct an incentive based on the ranking of the agents, so that a competition among the latter is initiated. We place ourselves in the limit setting of mean-field type interactions and prove the existence and uniqueness of the equilibrium distribution for a given reward, for which we can find an explicit representation. Focusing first on the homogeneous setting, we characterize the optimal reward function using a convex reformulation of the problem and provide an interpretation of its behaviour. We then show that this characterization still holds for a sub-class of heterogeneous populations. For the general case, we propose a convergent numerical method which fully exploits the characterization of the mean-field equilibrium. We develop a case study related to the French market of Energy Saving Certificates based on the use of realistic data, which shows that the ranking system allows to achieve the sobriety target imposed by the European commission.

Intermittency models based on integrated Volterra processes

Mireille Bossy (1),
Kerlyns Martínez (2),
Paul Maurer (1),

(1) Université Côte d'Azur, Inria, CNRS, France
(2) Universidad de Concepción, Chile.

Motivated by applications in physics (e.g., turbulence intermittency) and financial mathematics (e.g., rough volatility), this talk consider a family of integrated stochastic Volterra processes characterized by a small Hurst parameter $H < 1/2$. We investigate the impact of kernel approximation on the integrated process by considering the resulting weak error. We quantify this error in terms of the L_1 norm of the difference between the two kernels, as well as the L_1 norm of the difference of the squares of these kernels. Our analysis is based on a path-dependent Feynman-Kac formula and the associated path-dependent partial differential equation (PPDE).

We apply this weak error estimation to further analyse some causal process propositions for intermittent turbulent transport, corresponding to the case where H goes to zero.

Long time behavior of mean field games with a common noise

Pierre Cardaliaguet (1),
Raphaël Maillet (1),
Wenbin Yan (1),

(1) Université Paris Dauphine, Place du Mréchal de Lattre de Tassigny, 75775 Paris Cedex - France

Mean-field games describe stochastic control problems with a large number of controllers. Here, we focus on models in which the controller dynamics are further perturbed by common noise: the problem then becomes a coupled system of SPDEs. We will discuss the long-time behavior of this system.

The Kuramoto Mean Field Game

René Carmona (1), Quentin Cormier (2), Mete Soner (1),

(1) Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ, 08540, USA
(2) Inria, CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

We study the classical Kuramoto model as an infinite horizon mean field game. We show that when the interaction parameter is below a sharp critical value, the system exhibits incoherence: the uniform distribution is stable. Above this threshold, a bifurcation occurs and the game develops self-organizing time stationary solutions. I will first review the results for homogeneous players where all oscillators have the same intrinsic frequency. Then I will discuss the case of an heterogeneous population. In the latter, each oscillator has its own intrinsic frequency. Depending on the distribution of the intrinsic frequencies, new interesting phenomenon appear.

This is a based on a current work with René Carmona and Mete Soner.

Statistical Learning of Value-at-Risk and Expected Shortfall

David Barrera (1), Stéphane Crépey (2), Emmanuel Gobet (3), Hoang-Dung Nguyen (2,4), Bouazza Saadeddine (2,5),

(1) Departamento de Matematicas, Universidad de los Andes, Colombia.
(2) Université Paris Cité, Laboratoire de Probabilités, Statistique et Modélisation (LPSM), CNRS UMR 8001.
(3) École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France.
(4) Natixis, Paris.
(5) CACIB, Montrouge.

We propose a non-asymptotic convergence analysis of a two-step approach to learn a conditional value-at-risk (VaR) and a conditional expected shortfall (ES) using Rademacher bounds, in a non-parametric setup allowing for heavy-tails on the financial loss. Our approach for the VaR is extended to the problem of learning at once multiple VaRs corresponding to different quantile levels. This results in efficient learning schemes based on neural network quantile and least-squares regressions. An a posteriori Monte Carlo procedure is introduced to estimate distances to the ground-truth VaR and ES. This is illustrated by numerical experiments in a Student-t toy model and a financial case study where the objective is to learn a dynamic initial margin.

Risk Aversion of Insider and Asymmetric Information

Albina Danilova (1),
Valentin Lizhdvoy (2),

(1) Department of Mathematics, London School of Economics, UK.
(2) Department of Mathematics, Higher School of Economics, Russia.

We analyse how the risk aversion of insider affects the equilibrium in insider trading model with dynamic information. In particular, we will consider a dynamic information Kyle-Back model under new assumptions: a) exponential utility preferences of the insider, b) non- Gaussianity of the signal, and c) price set by the market maker being a function of weighted signal which is not necessarily Gaussian either. We will discuss conditions on the weighting and pricing functions which ensure the existence of equilibrium and derive, under afore mentioned conditions, the equilibrium pricing and weighting functions, as well as insider's optimal trading strategy.

Cooperation, competition, and common pool resources in mean field games

Gökçe Dayanıklı (1),
Mathieu Laurière (2),

(1) Department of Statistics, University of Illinois at Urbana-Champaign.
(2) Shanghai Frontiers Science Center of Artificial Intelligence and Deep Learning; NYU-ECNU Institute of Mathematical Sciences, New York University Shanghai.

The tragedy of the commons (TOTC, introduced by Hardin, 1968) states that the individual incentives will result in overusing common pool resources which in turn may have detrimental future consequences that affect everyone involved negatively. However, in many real-life situations this does not happen and researchers such as the Nobel laureate Elinor Ostrom suggested mutual restraint by individuals can be the preventing factor. In mean field games (MFGs), since individuals are insignificant and fully non-cooperative, the TOTC is inevitable. This shows that MFG models should incorporate a mixture of selfishness and altruism to capture real-life situations that include common pool resources. Motivated by this, we will discuss different equilibrium notions to capture the mixture of cooperative and non-cooperative behavior in the population. First, we will introduce mixed individual MFGs and mixed population MFGs where we also include the common pool resources. The former captures altruistic tendencies at the individual level and the latter models a population that is a mixture of fully cooperative and non-cooperative individuals. For both cases, we will briefly discuss definitions and characterization of equilibrium with the forward backward stochastic differential equations. Later, we will discuss a real-life inspired example of fishers where the fish stock is the common pool resource. We will analyze the existence and uniqueness results, and discuss the experimental results.

Mean-field control with Dirichlet-Fergusson common noise

François Delarue (1),
Mattia Martini (1),
Giacomo Sodini (2),

(1) Université Côte d'Azur, Nice, France
(2) Universität Wien, Austria

The aim of this presentation is to explore the regularizing effects of an infinite-dimensional common noise on mean-field control models. Ideally, it is expected that an infinite-dimensional common noise can enforce the uniqueness of solutions. However, the construction of such a forcing involves the introduction of a diffusion process taking values in the space of probability measures. Here, we study the impact of a Dirichlet-Fergusson type noise on a mean-field control problem and discuss the associated second-order Hamilton-Jacobi-Bellman equation.

Some Financial Applications of the Functional Itô Calculus

Bruno Dupire (1),

(1) Bloomberg

Path dependence is ubiquitous in finance, sometimes explicitly as the payoff of an exotic option may depend on the whole path of the asset price, not only at maturity, other times through the dynamics of the underlying (volatility, dividends. . .). The framework to model path dependence is the Functional Itô Calculus and we review its basic concepts before offering a partial panorama of its applications: computation of the Greeks of path dependent options, perturbation analysis, volatility risk decomposition, Taylor expansion with signatures for fast computation of VaR and characterization of attainable claims, amongst other ones.

An example of stochastic games arising in stock markets

Valdo Durrleman (1),

(1) Eisler Capital

An example of stochastic games arising in stock markets.

Rough stochastic differential equations

Peter Friz (1),

(1) Institut für Mathematik, Technische Universität Berlin, and Weierstrass Institute for Applied Analysis and Stochastics (WIAS), Berlin, Germany

Rough stochastic differential equations (RSDEs) unify Itô SDEs and Lyons' rough differential equations, offering a robust pathwise framework for stochastic modeling. Introduced by F-Hocquet-Lê in arXiv:2106.10340, RSDEs have since become a versatile tool across several domains, including

(I) pathwise stochastic filtering,

(II) pathwise control theory,

(III) local stochastic volatility modeling, and

(IV) conditional analysis of McKean–Vlasov equations under common noise.

This talk will survey these developments and highlight key mathematical structures.

Continuous-time mean field games: a primal-dual characterization

Xin Guo (1),
Anran Hu (1),
Jiacheng Zhang (2),
Yufei Zhang (3),

(1) UC Berkeley, IEOR, USA

In this talk, we will establish a primal-dual formulation for continuous-time mean field games (MFGs) and provide a complete analytical characterization of the set of all Nash equilibria (NEs). We first show that for any given mean field flow, the representative player's control problem with *measurable coefficients* is equivalent to a linear program over the space of occupation measures. We then establish the dual formulation of this linear program as a maximization problem over smooth subsolutions of the associated Hamilton-Jacobi-Bellman (HJB) equation, which plays a fundamental role in characterizing NEs of MFGs. Finally, a complete characterization of *all NEs for MFGs* is established by the strong duality between the linear program and its dual problem. This strong duality is obtained by studying the solvability of the dual problem, and in particular through analyzing the regularity of the associated HJB equation.

Compared with existing approaches for MFGs, the primal-dual formulation and its NE characterization do not require the convexity of the associated Hamiltonian or the uniqueness of its optimizer, and remain applicable when the HJB equation lacks classical or even continuous solutions.

A new approach to principal-agent problems with volatility control

Alessandro Chiusolo (1), Emma Hubert (1),

(1) Operations Research and Financial Engineering (ORFE), Princeton University.
Research partially supported by the NSF grant DMS-2307736.

The seminal work of Cvitanić, Possamaï and Touzi (2018) [1] introduced a general framework for continuous-time principal-agent problems using dynamic programming and second-order backward stochastic differential equations (2BSDEs). In this talk, we first propose an alternative formulation of the principal-agent problem that allows for a more direct resolution using standard BSDEs alone. Our approach is motivated by a key observation in [1]: when the principal observes the output process X continuously, she can compute its quadratic variation pathwise. While this information is incorporated into the contract in [1], we consider here a reformulation where the principal directly controls this process in a 'first-best' setting. The resolution of this alternative problem follows the methodology known as Sannikov's trick in continuous-time principal-agent problems. We then demonstrate that the solution to this 'first-best' formulation coincides with the original problem's solution. More specifically, leveraging the contract form introduced in [1], we establish that the 'first-best' outcome can be attained even when the principal lacks direct control over the quadratic variation. Crucially, our approach does not require the use of 2BSDEs to prove contract optimality, as optimality naturally follows from achieving the 'first-best' scenario. We believe that this reformulation offers a more accessible approach to solving continuous-time principal-agent problems with volatility control, facilitating broader dissemination across various fields. In the second part of the talk, we will explore how this methodology extends to more complex settings, particularly multi-agent frameworks.

Sharp quantitative propagation of chaos for mean field and non-exchangeable systems

Daniel Lacker (1),

(1) Columbia University

The propagation of chaos phenomenon states roughly that a large system of weakly interacting particles will remain approximately independent for all times if initialized as such. This can be quantified in terms of the distance between low-dimensional marginal distributions and suitably chosen product measures. This talk will discuss some recent sharp quantitative results of this nature, both for classical mean field diffusions and for more recently studied non-exchangeable models. These results are driven by a new "local" relative entropy method, in which low-dimensional marginals are estimated iteratively by adding one coordinate at a time, leading to surprising improvements on prior results obtained by "global" arguments such as subadditivity inequalities. In the non-exchangeable setting, we exploit a surprising connection with first-passage percolation.

An Efficient On-Policy Deep Learning Framework for Stochastic Optimal Control

Mathieu Lauriere (1),
Mengjian Hua (2),
Eric Vanden-Eijnden (3),

(1) Shanghai Frontiers Science Center of AI and DL; NYU-ECNU Institute of Mathematical Sciences, NYU Shanghai; Shanghai, China 200124.
(2) NYU-ECNU Institute of Mathematical Sciences, NYU Shanghai; Shanghai, China 200124
(3) Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

We present a novel on-policy algorithm for solving stochastic optimal control (SOC) problems. By leveraging the Girsanov theorem, our method directly computes on-policy gradients of the SOC objective without expensive backpropagation through stochastic differential equations or adjoint problem solutions. This approach significantly accelerates the optimization of neural network control policies while scaling efficiently to high-dimensional problems and long time horizons. We evaluate our method on classical SOC benchmarks as well as applications to sampling from unnormalized distributions via Schrödinger-Föllmer processes and fine-tuning pre-trained diffusion models. Experimental results demonstrate substantial improvements in both computational speed and memory efficiency compared to existing approaches.

Analysis in high dimension

Pierre-Louis Lions (1),

(1) Collège de France

In this talk we investigate limits, as the dimension goes to infinity, of point configurations, differential equations both deterministic and stochastic, functions and PDE's, transformations with or without permutation invariances.

Natural gas forward curves: from convenience yields to additive Gaussian processes

Mike Ludkovski (1),

(1) Department of Statistics & Applied Probability, University of California Santa Barbara

This is an empirical talk discussing the term structure of natural gas forward curves, which feature a rich structure across 100+ maturities. After reviewing extant literature, including the early work Carmona & Ludkovski (AMS Contemporary Mathematics, 2004), I will introduce a novel nonparametric decomposition of the forward curve into several additive components based on Gaussian process (GP) kernels. This decomposition includes periodic components to capture the annual seasonality, linear trends, and both local and global fluctuations, offering a new approach to holistically capture the entire shape in a single consistent fitting procedure. The probabilistic structure of the GPs furthermore lends itself to both extrapolation and infilling tasks in order to quantify uncertainty around illiquid forward quotes and very long maturities. I will conclude with a study of the best-fitting kernels, as well as empirical analysis based on a dataset comprising 15 years of history for CME Henry Hub and ICE TTF contracts. The presented results provide new insights into historical evolution and statistical dimension reduction of natural gas forward curves.

Applications of rough path theory

Terry Lyons (1),

(1) Mathematical Institute, University of Oxford

We survey applications of rough path theory.

A social discrete choice model under congestion

Roland Malhamé (1),
Jérôme Le Ny (1),
Noureddine Toumi (1),

(1) GERAD et École Polytechnique de Montréal

We consider a large population discrete choice model under congestion. Among other potential applications, this could capture crowd evacuation situations where an exit door has to be assigned to each agent while avoiding excessive congestion. For tractability reasons, we adopt a linear quadratic modeling framework. We simulate congestion effects by encouraging agents to spread maximally in the state space through a negative component in their individual cost function. Because of the non- convex nature of the problem, person by person optimization is excluded. Instead, we build a limiting large-scale optimal solution by considering first the exact solution of the finite population problem and subsequently sending the population size to infinity. This limiting solution is decentralized, more easily computable, and is shown to be ϵ-optimal. Numerical results are presented.

Second order mean-curvature flow as a mean-field game

Sergey Nadtochiy (1),
Yucheng Guo (2),
Mykhaylo Shkolnikov (3),

(1) Department of Applied Mathematics, Illinois Institute of Technology, Chicago, IL, USA.
(2) ORFE, Princeton University, Princeton, NJ, USA.
(3) Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA, USA.

The celebrated mean-curvature flow describes the evolution of the interface between two domains which moves so that its orthogonal velocity at each point is proportional to its mean curvature, pointing in the direction of decreasing the curvature. In the second order mean-curvature flow, it is the derivative of the orthogonal velocity (i.e., the acceleration) that is proportional to the mean curvature. Both flows can be described through partial differential equations (PDEs) for the associated arrival time functions. However, unlike the PDE for the classical mean-curvature flow, the equation for its second order version – which we refer to as the "cascade equation" – is hyperbolic and does not enjoy the comparison principle. For this reason, and due to other challenges, the standard PDE tools are not sufficient to develop a well-posedness theory for the cascade equation directly. Nevertheless, it turns out that solutions to the cascade PDE can be identified with minimal elements of a set of value functions in a family of mean-field games. As a result, the existence of a solution to the cascade equation can be shown by proving the compactness of the aforementioned set of value functions, which we accomplish by employing the tools from Geometric Measure Theory.

An inverse function theorem for Ito maps, with application to statistical inference for random rough differential equations

Anastasia Papavasiliou (1),
Thomas Morrish (1),

(1) Department of Statistics, University of Warwick, UK

Our goal is to develop a general framework for performing statistical inference for discretely observed random rough differential equations. The first step in our approach is to solve the 'inverse problem' under the assumption of continuous observations, i.e. construct a geometric p-rough path X whose response Y , when driving a rough differential equation, matches the observed trajectory y. We call this the continuous inverse problem and start by rigorously defining its solution. We then develop a framework where the solution can be constructed as a limit of solutions to appropriately designed discrete inverse problems, so that convergence holds in p-variation. Our approach is based on calibrating the rough path approximations whose limit defines the rough path 'lift' of X to the observed trajectory y. We prove the results for the case where X can be constructed as a limit of piecewise linear paths and discuss how this framework can be used to perform 'exact simulation' in-between observations.

Some Continuum Aggregate Games

Ronnie Sircar (1),

(1) ORFE Department, Princeton University.

Some recent game-theoretic models have interaction between players through the aggregate sum of their actions and/or states. This includes two models of cryptocurrency mining, Cournot-type models of oil production and prices, and two related models of renewable capacity expansion. We discuss an approach in which a sum is approximated by the (large) number of players multiplied by the continuum mean, in order that mean field game analysis can be applied. We study convergence of equilibria in very simple (static) settings and give preliminary general results.

A statistical test to show evidence of the infinite expectation of René's energy

Denis Talay (1),
Hector Olivero (2),

(1) Inria Saclay
(2) Universidad de Valparaiso

Motivated by the simulation of stochastic particle systems with singular McKean-Vlasov interaction kernels and by observations of René's energy expenditure process we address the problem of detecting that a sampled probability distribution has infinite expectation.

More precisely, we construct and analyze an hypothesis test whose null hypothesis is: 'The sampled distribution is in the domain of attraction of the Normal law' and the alternative hypothesis is: 'The sampled distribution is in the domain of attraction of a stable law with index smaller than 2'. Our key observation is that the sampled distribution cannot have a finite second moment when the null hypothesis is rejected.

Surprisingly, we find it useful to derive our test from the statistics of semimartingales: More precisely, from tests aimed to determine whether a semimartingale has jumps by observing one single path at discrete times.

We discuss the choice of crucial parameters involved in the test and illustrate our theoretical results with numerical experiments.

Propagation of carbon price shocks through the value chain: the mean-field game of defaults

Zorana Grbac (1),
Simone Pavarana (2),
Thorsten Schmidt (2),
Peter Tankov (3),

(1) LPSM, Université Paris Cité
(2) University of Freiburg
(3) CREST, ENSAE, Institut Polytechnique de Paris

We develop a mean-field game framework to assess the impact of carbon pricing in a multisectoral economy composed of multiple defaultable firms in each sector. In our model, each sector produces a homogeneous good, the price of which is endogenously determined through market clearing. Firms act as price takers and maximize their profits by optimally combining inputs — including intermediate goods from other sectors, emissions, and labour — while interacting with other firms via the sectoral good's price. Firms select optimal default dates to maximize shareholder value.

Taking the mean-field limit, we reformulate the economy as an optimal stopping mean-field game within each sector. We then solve the resulting system of coupled mean-field games using the linear programming approach, which characterizes Nash equilibria in terms of population measure flows. We prove the existence of a linear programming Nash equilibrium and establish the uniqueness of its corresponding price system.

Numerical illustrations are provided for firms with CES production functions in a stylized two-sector economy (green and brown sectors). Our experiments demonstrate that carbon price shocks can trigger significant spillover effects across sectors, highlighting the importance of sectoral interdependencies in shaping decarbonization pathways.

The dynamics of martingale marginals

Michael Tehranchi (1),

(1) Statistical Laboratory, University of Cambridge, United Kingdom

This talk revisits the problem of modelling the dynamics of the set of conditional marginal laws of a martingale. It will survey the work of Carmona and others on the Heath–Jarrow–Morton-type evolution for the implied volatility surface, as well as some newer ideas.

Cortical wavelets and sparse Bayesian learning for brain source localization and reconstruction

Bruno Torrésani (1),
Samy Mokhtari (1,2),
Jean-Michel Badier (2),
Christian Bénar (2),

(1) Institut de Mathématiques de Marseille (I2M), Aix-Marseille Université, CNRS
(2) Institut de Neuroscience des Systèmes (INS), Aix-Marseille Université, INSERM.

Electro-encephalography and magneto-encephalography are non-invasive imaging modalities that provide measurements of time-dependent measurements of brain activity with high temporal accuracy and good spatial accuracy. However, the number of sensors is generally much smaller than the number of unknowns, which makes the corresponding inverse problem extremely ill conditioned.

Accurate solutions can be obtained when the unknown brain activity can be represented sparsely in some specific domain. This is often done by minimizing a suitable cost function, that involves a non-smooth, parameter-dependent regularization. The parameter tuning is itself a difficult issue.

This work focuses on extended (i.e. non-focal) brain activity, and exploits sparsity in a spatial wavelet domain. Here, spatial wavelets means wavelet systems defined on the surface of the subject's cortex, following a construction by Hammond, Gribonval and Vandergheynst based upon graph Laplacian eigen-decomposition. The parameter tuning problem is circumvented using the so-called sparse Bayesian learning (SBL) approaches, which provide sparse, diagonal estimates for hyperparameters of the prior distribution of unknowns. This results in estimates for brain activity that are sparse in the wavelet domain, and extended in the spatial domain.

The talk will present the wavelet-SBL approach, and numerical results on simulated and real data.

Asymptotics of Yule's nonsense correlation: can one test the dependence of two random walks as they converge in law to Brownian paths?

Frederi Viens (1),
Philip Ernst (2),
John (Tripp) Roberts (1),
Shuo Yan (2),

(1) Department of Statistics, Rice University, Houston, TX, USA. (2) Department of Mathematics, Imperial College London, UK.

In 1926 Statistician G. Udny Yule showed that for two independent standard random walks, the empirical correlation coefficient (Pearson's correlation) does not converge to 0, but rather appears to converge in distribution to a diffuse law supported by the entire interval (-1,1). This phenomenon, which has since been recognized for many highly non-stationary time series, is in sharp contrast with the classical result for two sequences of i.i.d. data, by which the same correlation converges to zero, a phenomenon which also extends to many stationary time series. Still, ignorance of this empirical fact for random walks and other non-stationary time series, known today as Yule's nonsense correlation, has lead practitioners to make dramatically ill-informed assertions about statistical associations. This improper use of methodology has occurred in recent times, particularly in environmental observational studies, e.g. for attribution in climate science. The mathematics behind the basic premise of Yule's nonsense correlation are a rather straightforward application of the classical Donsker's theorem; the Pearson correlation ρ_n of two random walks of length n converges in distribution to the law of a random variable ρ written explicitly as the ratio of two quadratic functionals of two Wiener processes on [0.1]. In this talk, we investigate the fluctuations around this convergence. We present elements of a new result by which $n(\rho - \rho_n)$ has an asymptotic distribution in the so-called second Wiener chaos, whose characteristics are partly exogenous to the original data, as one would expect for a standard central limit theorem, and are partly conditional on the data. We will discuss the implications of this discovery in practical testing for independence and for attribution in environmental time series. We conjecture that the fluctuation scale, of order $1/n$ rather than $1/n^{1/2}$, is not accidentally related to the exotic convergence in law in the second Wiener chaos.