

Abstracts of the talks at GTP 2014

CIMAT, Guanajuato, Mexico

13–15 November, 2014

Jacob Abernethy *Information aggregation in exponential family markets*
(joint work with Sindhu Kutty, Sebastien Lahaie, and Rahul Sami)

We consider the design of prediction market mechanisms known as automated market makers. We show that we can design these mechanisms via the mold of exponential family distributions, a popular and well-studied probability distribution template used in statistics. We give a full development of this relationship and explore a range of benefits. We draw connections between the information aggregation of market prices and the belief aggregation of learning agents that rely on exponential family distributions. We develop a very natural analysis of the market behavior as well as the price equilibrium under the assumption that the traders exhibit risk aversion according to exponential utility. We also consider similar aspects under alternative models, such as when traders are budget constrained.

Dmitry Adamskiy *A closer look at adaptive regret*

For the prediction with expert advice setting, we consider methods to construct algorithms that have low adaptive regret. The adaptive regret of an algorithm on a time interval $[t_1, t_2]$ is the loss of the algorithm minus the loss of the best expert over that interval. Adaptive regret measures how well the algorithm approximates the best expert locally, and so is different from, although closely related to, both the classical regret, measured over an initial time interval $[1, t]$, and the tracking regret, where the algorithm is compared to a good sequence of experts over $[1, t]$.

We investigate two existing intuitive methods for deriving algorithms with low adaptive regret, one based on specialist experts and the other based on restarts. Quite surprisingly, we show that both methods lead to the same algorithm, namely Fixed Share, which is known for its tracking regret. We provide a thorough analysis of the adaptive regret of Fixed Share. We obtain the exact worst-case adaptive regret for Fixed Share, from which the

classical tracking bounds follow. We prove that Fixed Share is optimal for adaptive regret: the worst-case adaptive regret of any algorithm is at least that of an instance of Fixed Share.

Elie Ayache *The market of contingent claims*

We ask how the market of contingent claims can be formalized or formally deduced. This is problematic because if we scrutinize the mathematical formalism of derivative pricing theory (or the pricing theory in contingent claims), for instance in a theoretical paper such as Harrison & Pliska (1980) and not in a heuristic paper such as the original Black & Scholes (1973), we find that there is only mention of the market of the underlying asset. With the underlying asset we can only manufacture or synthesize contingent payoffs, not contingent claims. The distinction is important. When the underlying is given and is trading with Brownian motion in its market, we may buy it or sell it and dynamically readjust our holding in order to end up with any payoff profile as a function of S , or $f(S)$, at any time horizon, aka contingent payoff. This is the major discovery of Black, Scholes and Merton. But the contingent payoff that we have manufactured is no contingent claim. Nobody has given us permission, in the formalism, to assume that contingent claims or derivatives independently exist as written contracts on which such contingent payoffs have been marked, and of which we subsequently wonder whether we can replicate them with the underlying.

In the formalism, when it is properly read, there is no mention of the contingent claim, or of an independent paper that may admit of an independent price, and of which we later argue that its price is curbed by the replication. In other words, the market of contingent claims doesn't exist in the formalism and only the market of the underlying does. So it is not that contingent claims are redundant (as it is traditionally thought) and for this reason they won't admit of their own market. Rather, they simply don't exist. Because they don't exist in the formalism, the formalism doesn't say they are redundant, hence it may become an impeccable consequence or conclusion of the formalism that they admit of a market. This is how we impeccably derive their market from the formalism, or rather, outside the formalism. The reason why we don't discard the formalism altogether and why we insist that the market of contingent claims should be its "consequence" is that the formalism is needed by the market-maker in order to hedge the contingent claims. There wouldn't have been an industry of contingent claims and their market wouldn't have been cemented without the historical layer of the formalism. Yet the matter of this market lies liter-

ally in the void outside the formalism. We are literally talking about an ineptual medium, properly material and not formal, properly historical, in which alone the market of contingent claims can be thought and given form, i.e., “formalized”, albeit meta-formally.

Jasper de Bock *Imprecise probability trees: a comparison of two different approaches and their limit behaviour*

I intend to compare two different ways of combining the local models that are attached to an event tree: the game-theoretic approach, and the behavioural approach (in the sense of Walley). I will discuss and illustrate the difference between the joint models that correspond to these two approaches, will briefly discuss the extent to which these differences are relevant in practice, and will present some new results on the limit behaviour of these two approaches.

Evgeny Burnaev *Non-parametric volatility estimation based on wavelet transform*

GARCH models, widely used in industry for volatility modeling and forecasting, are linear, parametric and stationary. However, real financial data are non-stationary and exhibits various non-linear phenomena. Therefore, non-parametric non-stationary model should be used. We model volatility by piece-wise constant non-parametric function and use wavelet transform based thresholding algorithm for its identification. The main difficulty here, opposed to standard wavelet denoising schemes, is to set up thresholds in a proper way due to non-homogeneity of wavelet coefficients variances. Experiments on real data showed efficiency of proposed approach.

Gert de Cooman *A game-theoretic ergodic theorem for imprecise Markov chains*

Abstract to appear.

Mark Davis *Pathwise local time and A prequential approach to financial risk management*

Abstracts to appear.

A. Philip Dawid *Fundamentals of Prequential Analysis*

Prequential (= Predictive Sequential) Analysis is a very general approach to problems of statistical inference and data analysis that takes seriously the

problem of making and criticising sequential forecasts. On the theoretical front, it allows for broad extension of classical statistical ideas such as estimation efficiency or consistent model selection, and has many pleasing properties. Practically, it provides a more generally applicable, interpretable and justifiable alternative to cross-validation. It has close links with Bayesian Inference, with Stochastic Complexity, with Algorithmic Complexity Theory, with Computational Learning Theory, with Rational Learning in Games, and with Game-Theoretic Probability. This expository presentation will give an overview of its main features.

Miroslav Dudik *A combinatorial prediction market for the US Elections* (joint work with Sebastien Lahaie, David Pennock, and David Rothschild)

Prediction markets are emerging as a powerful and accurate method of aggregating information from populations of experts (and non-experts). Traders in prediction markets are incentivized to reveal their information through buying and selling securities for events such as “Republicans will win the US Senate majority in 2014”. The prices of securities reflect the aggregate belief about the events and the key challenge is to correctly price the securities.

I will present a new algorithm (an automated market maker) for pricing multiple logically interrelated securities. Our approach lies somewhere between the industry standard—treating related securities as independent and thus not transmitting any information from one security to another—and a full combinatorial market maker for which pricing is computationally intractable. Our techniques borrow heavily from variational inference in exponential families. We prove several favorable properties of our scheme and evaluate its information aggregation performance on survey data involving hundreds of thousands of complex predictions about the 2008 US presidential election. I will also give an example of the real-world deployment of our pricing scheme in the 2012 US presidential elections where we modeled probabilities of 10^{33} related events and updated them in real time.

Rafael M. Frongillo *Possibility and impossibility of liquidity adaptation in prediction markets*

We study two regimes in the design of cost-function-based combinatorial prediction markets, one where the designer wishes liquidity to increase as the market progresses, and the other where the market maker would like liquidity with respect to certain information to decrease.

For increasing liquidity, we introduce a framework for automated market

making for prediction markets, the volume parameterized market (VPM), in which securities are priced based on the market maker’s current liabilities as well as the total volume of trade in the market. We show that many existing market makers fall into this framework as special cases, and design a new market maker, the perspective market, that satisfies four out of five desirable properties in the complex market setting, but fails to satisfy information incorporation. However, we show that the sacrifice of information incorporation is unavoidable: we prove an impossibility result showing that any market maker that prices securities based only on the trade history cannot satisfy all five properties simultaneously.

For decreasing liquidity, we assume that it is known that some piece of information will be revealed to traders, and the market maker wishes to prevent guaranteed profits for trading on the sure information. We design adaptive cost functions which: (1) preserve the information previously gathered in the market; (2) eliminate (or diminish) rewards to traders for the publicly revealed information; (3) leave the reward structure unaffected for other information; and (4) maintain the market maker’s worst-case loss.

Yuri Gurevich *Impugning alleged randomness*

According to a 1985 issue of New York Times, “The New Jersey Supreme Court today caught up with the Essex County Clerk and a Democrat who has conducted drawings for decades that have given Democrats the top ballot line in the county 40 times out of 41 times.” In the 1980s the Israeli tax authorities encouraged the public to request invoices (from plumbers, painters, etc.) and send the invoices in; big prices were ruffled off. But the operation collapsed when it turned out that the winner was none other than the Director of Customs and VAT at the time.

You may be convinced that such lotteries are rigged, but how would you justify your assertion in the court of law? We attempt to furnish you with an argument.

Yuri Kalnishkan *Practical prediction with specialist experts*

Prediction with expert advice is concerned with mixing predictions made by pools of “experts”, which can be prediction algorithms, humans, or other kinds of predictors treated as black boxes. A method of prediction with expert advice comes with performance guarantees ensuring that the learner employing this method will do nearly as well as the best expert whatever the circumstances.

Specialists (or sleeping experts) are experts that can refrain from making

a prediction (sleep) on certain steps. A performance guarantee for the master learning algorithm should ensure that it performs on a par with every specialist expert from the pool on the steps where the specialist expert does not sleep. Specialist experts were introduced by Y. Freund et al. in 1997, but the methods for merging them were rather cumbersome until simplified by A. Chernov and V. Vovk in 2009.

This talk will discuss applications of specialist experts to the following problem. When attempting to predict certain events, one is often faced with the task of selecting relevant historical information. Sometimes only the most recent information matters; this leads to sliding window or decaying weights techniques. However, sometimes old information becomes relevant again. For example, at a time of economic recession one may want to look for the previous recession period rather than more recent period of steady growth. Simple specialist experts can often capture different forms of dependency on the past and merging techniques from prediction with expert advice automatically achieve optimal trade-offs among them.

Several examples of datasets, such as implied volatility of options and students' performance in tests, and specialist experts techniques for them will be discussed.

Wouter Koolen *Buy low, sell high*

We consider online trading in a single security with the objective of getting rich when its price ever exhibits a large upcrossing, without risking bankruptcy. We investigate payoff guarantees that are expressed in terms of the extremity of the upcrossings in the observed price sequence. We obtain an exact and elegant characterisation of the guarantees that can be achieved, and explore some delightfully intricate example trading strategies. Moreover, we derive a simple canonical strategy for each attainable guarantee.

Kathryn Laskey *SciCast: A combinatorial prediction market for science and technology forecasting*

Abstract to appear.

Robert Meister *Just Optionality* (TBC)

Abstract to appear.

Kenshi Miyabe *Derandomization in game-theoretic probability* (joint work with Akimichi Takemura)

I will explain a general method for constructing a concrete deterministic strategy of Reality from a randomized strategy in game-theoretic probability. The construction can be seen as derandomization in game-theoretic probability. I also introduce probability notions relating Reality's strategy while usual probability notions in game-theoretic probability is related Skeptic's strategy.

Shota Nakagawa *Games for discrete-time Markov chain and their application to verification*

We formulate some notions and results about discrete-time Markov chains in terms of game-theoretic probability. Discrete-time Markov chain is one of the models which are used to model systems' probabilistic behavior in formal verification. We demonstrate that some notions and results like fairness theorem and simulations are natural and simple in terms of game-theoretic probability.

Rimas Norvaisa *Power variation and p -variation of sample functions of stochastic processes*

Abstract to appear.

Ilia Nouretdinov *Power variation and variation index* (joint work with Vladimir Vovk)

This talk will discuss the relationship between power variation (including the standard notions of total variation and quadratic variation as special cases) and variation index, which are often used as a measure of volatility of a function. Both measures of volatility are defined for a given sequence of partitions. We consider two natural classes of sequences of partitions, Riemann and Lebesgue dyadic sequences; whereas for the latter the relationship is very close, simple examples show that for the former it is extremely loose.

Harald Oberhauser *Some applications of the signature*

The theory of rough paths has led to new insights on how to understand the evolution of systems that evolve under randomness, e.g., stochastic (partial) differential equations. A central tool is the so-called signature of a path. In this talk I will report about some recent work that combines some ideas from rough paths with machine learning techniques.

Alexander V. Outkin *On applications of game-theoretic probability and defensive forecasting to agent-based models of markets*

We present an attempt on connecting agent-based modeling with game-theoretic probability and defensive forecasting and outline a framework connecting elements of game-theoretic probability with agent-based models. We illustrate this framework on an example of our model of the Nasdaq stock market and on an example of a model of a natural gas market model, and show how game-theoretic probability can be used to test the simulated market price dynamics, the individual agent trading strategies, and the overall agent-based model.

Dusko Pavlovic *Learning equilibria in randomized monoidal computation*

Probability is usually interpreted either as a result of an objective process of counting and frequency calculations, or as an expression of subjective beliefs. In a related dichotomy, the market is usually interpreted either as an objectively randomized process of equilibrium selection, or as the realm of subjective bubbles that keep blowing and bursting and blowing again. Either way, the players on the market demonstrate their unbounded rationality by using probability theory. The other way around, probability theory can be instructively derived from the market players' rational behaviors.

But as markets have migrated into networks of computers, their behaviors have become both more objective, in the sense that they are harder to “program” or even regulate, and more subjective, in the sense that the bubbles and delusions grow faster, and take more complex forms. Can we still get away without taking computation into account in our probability, decision and game theories?

In this talk, I will report about some preliminary results in overcoming the well-known obstacles to computational game theory and bounded rationality. The main tool of this research is monoidal computer [Information and Computation 226, arXiv:1208.5205, arXiv:1402.5687], a graphic model of computation that evolved from attempts to avoid “machine programming” in computability and complexity theoretic proofs and constructions. I will sketch a basic feasible equilibrium construction, and explain the suitable form and role of randomization in this framework.

Glenn Shafer *Introduction to the workshop and game-theoretic probability*

No abstract.

Pietro Siorpaes *Pathwise versions of the Burkholder-Davis-Gundy inequality*

We present a new proof of the Burkholder-Davis-Gundy inequalities for $1 \leq p < \infty$. The novelty of our method is that these martingale inequalities are obtained as consequences of elementary deterministic counterparts. The latter have a natural interpretation in terms of robust hedging.

Akimichi Takemura *A game-theoretic proof of the Erdős-Feller-Kolmogorov-Petrovsky law of the iterated logarithm for fair-coin tossing* (joint work with Takeyuki Sasai and Kenshi Miyabe)

We give a game-theoretic proof of the celebrated Erdős-Feller-Kolmogorov-Petrovsky law of the iterated logarithm for fair coin tossing. Our proof, based on Bayesian strategy, is explicit as many other game-theoretic proofs of the laws in probability theory.

Ambuj Tewari *Sequential complexities and uniform martingale laws of large numbers* (joint work with Alexander Rakhlin and Karthik Sridharan)

I will describe necessary and sufficient conditions for a uniform martingale Law of Large Numbers. These conditions are obtained by extending the technique of symmetrization from iid random variables to the case of dependent random variables. Such an extension provides “sequential” or “online” analogues of various classical measures of complexity, such as covering numbers, combinatorial dimensions, and Rademacher averages from empirical process theory.

Vladimir Vovk *Laws of probabilities in efficient markets*

In this talk I will argue that different kinds of markets force different properties of price paths. Only a few examples will be given; the general picture is still missing at this time.

Vladimir V'yugin *Log-optimal portfolio selection using method of calibration*

We study sequential investment strategies for financial markets. Investment strategies are allowed to use information collected from the past of the market and determine, at the beginning of a trading period, a portfolio, that is, a way to distribute their current capital among the available assets. The goal of the investor is to maximize his wealth in the long run using minimal assumptions on market mechanisms.

In modeling the behavior of the evolution of the market, two main approaches have been considered in the theory of sequential investment. In the first of them we assume that the market vectors are realizations of a random process, and describe a statistical model. Another “worst-case” approach allows the market sequence to take completely arbitrary values, and no stochastic model is imposed on the mechanism generating the price relatives.

Following “worst-case” approach, we present a method for log-optimal portfolio selection based on “an artificial” probability distribution of market sequence constructed by the method of Foster–Vohra calibration. Our portfolio is log-optimal with respect to the class of all continuous investment strategies using side information.