Regrets, learning and wisdom

Damien Challet

► To cite this version:

Damien Challet. Regrets, learning and wisdom. The European Physical Journal. Special Topics, 2016, 225 (17-18), pp.3137 - 3143. 10.1140/epjst/e2016-60122-y . hal-01312973

HAL Id: hal-01312973

<https://hal-centralesupelec.archives-ouvertes.fr/hal-01312973>

Submitted on 21 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Regrets, learning and wisdom

Damien Challet

Laboratoire de Mathématiques et Informatique
pour la Complexité et les Systèmes
CentraleSupélec, Université Paris-Saclay, France
Encelade Capital SA, Lausanne, Switzerland
`damien.challet@centralesupelec.fr`

October 2, 2018

Abstract

This contribution discusses in what respect Econophysics may be able to contribute to the rebuilding of economics theory. It focuses on aggregation, individual vs collective learning and functional wisdom of the crowds.

1 Introduction

A good starting point to rebuild an economic theory is to use agent-based models that include learning, interaction and networks [29, 7, 8, 4]. This framework is a natural meeting point for Economics and Physics (and Psychology and Biology and Computer Science and ...), which already hints that Econophysics is only part of the solution.

Statistical Physics's strength comes from its familiarity with collective phenomena. Aggregating the non-linear actions of many interacting individuals leads to remarkable global phenomena and great mathematical simplifications [36, 39]. Whether the outcome is optimal for the agents or the system is of central importance. This contribution argues that learning and optimality may occur at various levels, may be either implicit or explicit, and that Econophysics would be wise to incorporate some more ideas from Neuroscience, Computer Science and Experimental Psychology.

Let me start with a few generic remarks about the difference of approaches to data analysis and modelling in Econophysics and Economics or Finance.

2 Interdisciplinary communication

The grass in other fields seems not only greener but also disconcerting at times. This, of course, works both ways.

Figure 1: Starlings flocking and a line. Original image source: Wikipedia.org

2.1 Statistics

In their 10-year old worries about Econophysics [21], Gallegati, Keen, Lux and Ormerod pinpointed the general disregard of econophysicists for Statistics, and rightly so. The spontaneous reaction of physicists is/was to assume to have collected enough data to dispense with statistical tests and tables. The situation has much changed in this respect. Confidence intervals of estimates are not unheard of nowadays. Physicists not only use some statistical tools, but even propose new statistics [44, 35, 12, 23].

2.2 Star(t)ling fits

The abundance of linear fits in Economics and Finance papers often puzzles physicists. Let me discuss what model fitting implies generically. Measuring something is equivalent to projecting a system into a sub-space. A good example is that of a picture taken by a camera: it is projection of a 3+1-dimensional world into a 2-dimensional world. In addition, the position of the camera is also of crucial importance. Figure 1 shows of a flock of bird. For centuries, people have puzzled about the 3-dimensional structure of such flocks, before it was realized that these clouds were dynamical two-dimensional objects, i.e., ribbons [10]. Had anyone been able to ask a bird what the shape was like from the inside, the shape of starlings' flocks would be have been known a long time ago. In short, placing oneself in the right space is a necessary condition for meaningful fits.

Taking linear regressions against a few well-known factors is the same thing as taking a picture from a limited number of popular standpoints. Say that one wishes to analyse the performance of a collection of hedge funds. A questionable approach is to project their performance in the space spanned by the few Fama-French factors [18] and then argue that hedge funds trade on such and such factors. This is inherently incomplete and very unlikely to yield real understanding of what drives hedge

funds performance. What one really needs to do is to reverse-engineer the performance of hedge funds. This requires to place one self in a space that encompasses all likely used trading strategies, or equivalently, to define factors as the returns of a variety of these strategies. Reference [46] is mostly successful in replicating the returns of several thousand funds, except two: the feeder fund of Madoff, which lived in a fantasy world, and Capital Fund Management Stratus Fund because the chosen strategy space did not include the strategies of that particular fund, which resulted in an meaningless projection.

Fitting any kind of model, however more sophisticated, to data is a projection. For example, calibrating a simple agent-based model [28] to financial data based on a moment method is a double projection, i.e., a double dimensionality reduction: from the data to moments and from the model to moments [1]. Model and data meet in a third space. This may yield an incomplete fit, which therefore may not be more efficient than other approaches. This leads us to microscopic models.

3 Lurking Ising models

Initially, econophysicists tried to apply the models they were most familiar with to financial or economic situations, which sometimes was hair-raising, even for some physicists. There is indeed no reason why financial markets should be exactly equivalent to a gas of electrons, even if some random power-law exponent seems right. This was deeply worrying. Much has changed since then.

The case of the Ising model is less controversial, if only because it is equivalent to Schelling's model [40] and because it is easy to see why it is quite likely to appear in discrete-choice agent-based models. It also illustrates the variety of what a classical spin may describe in other contexts.¹ The simplest idea of course is to map the two possible values of a classical spin to two opposite decisions, which seems natural for investors [14, 25, 6]. Another possibility is to map two alternative possibilities to the two spin values: Ref. [37] proposes a test for the presence of social imitation in the choice between two alternative possibilities. The key point is that this test is based on exact results from mean-field random-field Ising model and consists in non-linear relationships between two quantities, all of which would be impossible to guess *a priori*. In other words, analytically tractable aggregation provides much more than moments of intellectual satisfaction.

Yet, all the above models are built as Ising models from scratch: actions are directly mapped to classical spins. Assuming instead that the binary choice is which strategy to use also leads to disordered spin models: agents with very limited possible actions in a complex world are able to optimise a global quantity, which can be written as a mean-field spin Hamiltonian [13] where the disorder comes from agents' heterogeneity. When binary choices are involved in interacting agent models, it is hard to avoid Ising models.

¹For a review of the Ising model in Econophysics, see [41] and references therein.

4 Learning

4.1 Microscopic learning

Whereas Logit learning is often found in Econophysics literature, surprisingly few other types of learning have been investigated. There lies much potential to establish bridges with other disciplines. First, computer scientists have also applied learning to finance [15]. More generically, in a Markovian context, Q-learning rests on the assumption that the system may be classified in a finite number of states by the agents and converges towards the optimal policy [45]. Computer scientists duly applied this scheme to the Minority Game, for example, defining a state as either which strategy was used, or the previous correct decision. Although this is not discussed in their papers, such dynamics seem to converge to a Nash equilibrium (e.g. [2]). Methods from Statistical Physics are without any doubt able to tackle this kind of learning scheme.

Another central question is what to learn. The current consensus in Neuroscience is that we learn to regret what we did not do and that a kind of Q-learning describes well how the brain works [38]. In the context of financial markets, this reinforces bubbles and crashes [33]. Indeed, when the price of an asset has an apparent trend, investors that do own any shares of the said asset regret not to have invested earlier, which triggers their investment. Reversely, when they are invested in an asset whose price begins to fall, they regret not to have closed their positions earlier.

Finally, Econophysicists have incorporated remarkably few well-known behavioural biases in their models, as they often assume that agents are risk-neutral. True, there are many reasons why one should use the exponential or logarithmic utility functions with about a ton of salt. However, it is surprising that even Prospect Theory [27] is nowhere to be seen in our agent-based models. By contrast, in their beautiful paper [3], Barberis *et al.* further simplify Prospect Theory with linear approximations, keeping the crucial feature that losses are about twice as painful as gains for agents, and a reference point which is a moving average of past wealth. Prospect Theory hence requires to distinguish between gains and losses in the agent payoff equation updates, which inevitably leads to additional complications.

Adding a pinch of Prospect theory in our agent-based models is doable. For example, asymmetric gains and losses can be included in the Minority Game and turn out to be another cause for the emergence of large fluctuations [5]. There is little doubt that De Dominicis generating functionals [17] can accomodate Prospect Theory in more complex agent-based models.

5 Systemic learning

The Darwinian force in financial markets that makes them adaptive systems has long been noted [19, 48, 32]. It does not imply however that the agents themselves are adaptive (e.g. that speculative funds calibrate their strategies in real-time). Indeed, at a global level (and at long time scales),

the relative importance of a single agent or of a sub-population of agents may evolve because of competitive processes (for food, wealth, price predictability, etc.): fitness selection is an indirect form of global learning (see e.g. [20]). As a result, even economic systems with zero-intelligence agents undergo a global learning process provided that some kind of selection is performed, as reflected by replicator equations in Evolutionary Game Theory [47]. This does not systematically result in a measurable global optimality.

There are cases, however, where a single global quantity is fairly well optimised, sometimes for obvious reasons (e.g. financial predictability). This is the Vox Populi [22], or Widsom of Crowds [42] effect: the aggregation of inconsistent opinions may lead to consistent aggregate estimates. Although well-known historical examples are really about estimating a single outcome such as the position of a lost craft or the weight of some object, the power of aggregation extends to much more generic situations.

Ensemble learning such as Random Forests [9] apply this idea to classification and regression problems. The latter kind of problem implies that an ensemble of imperfect learners may correctly learn functional relationships, a much more difficult task. The implication for Economics is that the many textbook noiseless “laws”, for example between price and excess demand, may be valid at an aggregate level. In other words, the clearly too simplistic economic intuition exposed in standard textbooks are in fact quite noisy relationships,. In passing the noise may be due in part to the heterogeneity of economic entities. A most striking illustration of how to extend these “theories” comes from a work on Marseilles Fish Market [24]: one of its figures plots the price paid for a type of fish as a function of the quantity sold for many transactions. A cloud of point emerges. Only when local averages are taken does emerge a relationship similar to those predicted by usual economic theories. More recently, collective portfolio optimisation in the presence of a complex transaction cost structure was found in brokerage data [16].

The conditions under which collective learning may occur are still unclear and provide a nice challenge for the years to come [34, 11] and certainly one which Physicists can contribute to. Finding more examples of wisdom of the crowds will also be part of the fun. Beyond the average behavior, the origin and role of heterogeneity in the dynamics of these systems are complementary research topics.

6 Conclusion

The question is not if Economics can become a physical science, but how to make it a science. 20 years of multidisciplinary research have convinced me that it will not be enough to sprinkle economic theory with a few more mechanistic ingredients and to take a slightly less axiomatic approach. A famous speech by Jean-Claude Trichet called for help from a wide range of hard and soft sciences [26]. Only the combination of Biology, Experimental Psychology, Computer Science and Physics is likely to make a difference. It is hard to disagree with this point of view [29, 7, 8, 41, 4].

Using common tools and concepts will certainly help achieving better

cooperation. Agent-based models, learning and networks certainly qualify as common grounds. A good example of cooperation between economists and physicists is the European CRISIS project, which lead to substantial scientific cross-fertilization. More sophisticated tools of Statistical Physics such as generating functionals are nowadays used by mathematical economists, which is great news, but, expectedly, only by the most mathematically minded, very much as in Physics. Mean-field games [31] will also contribute to establish bridges between Physics and Economics (see e.g. [43]).

With regards to whether less open-minded economists will accept the resulting new economic thinking, there are many reasons to be optimistic. Each economic crisis is an Economics crisis [29], and leads to more realistic models. For example, the 2008 crisis has triggered much interest in real networks and self-excited processes (e.g. [30]).

Let us build it and they will use it.

References

- [1] Simone Alfarano, Thomas Lux, and Friedrich Wagner. Estimation of a simple agent-based model of financial markets: An application to australian stock and foreign exchange data. *Physica A: Statistical Mechanics and its Applications*, 370(1):38–42, 2006.
- [2] M Andrecut and MK Ali. Q learning in the minority game. *Physical Review E*, 64(6):067103, 2001.
- [3] Nicholas Barberis, Ming Huang, and Tano Santos. Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53, 2001.
- [4] Stefano Battiston, J Doyne Farmer, Andreas Flache, Diego Garlaschelli, Andrew G Haldane, Hans Heesterbeek, Cars Hommes, Carlo Jaeger, Robert May, and Marten Scheffer. Complexity theory and financial regulation. *Science*, 351(6275):818–819, 2016.
- [5] Federico Bettin. The role of asymmetric gain and loss perception in minority games. Master’s thesis, Università degli Studi di Padova, 2014.
- [6] Stefan Bornholdt. Expectation bubbles in a spin model of markets: Intermittency from frustration across scales. *International Journal of Modern Physics C*, 12(05):667–674, 2001.
- [7] Jean-Philippe Bouchaud. Economics needs a scientific revolution. *Nature*, 455(7217):1181–1181, 2008.
- [8] Jean-Philippe Bouchaud. The (unfortunate) complexity of the economy. *Physics World*, 22(04):28, 2009.
- [9] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [10] Andrea Cavagna, Alessio Cimorelli, Irene Giardina, Giorgio Parisi, Raffaele Santagati, Fabio Stefanini, and Massimiliano Viale. Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences*, 107(26):11865–11870, 2010.

- [11] L Elisa Celis, Peter M Krafft, and Nathan Kobe. Sequential voting promotes collective discovery in social recommendation systems. *arXiv preprint arXiv:1603.04466*, 2016.
- [12] Damien Challet. A robust, efficient, and moment-free estimator of Sharpe ratios. Working paper, available at <http://arxiv.org/abs/1505.01333>, 2015.
- [13] Damien Challet, Matteo Marsili, and Yi-Cheng Zhang. *Minority Games*. Oxford University Press, Oxford, 2005.
- [14] R. Cont and J.-P. Bouchaud. Herd behaviour and aggregate fluctuation in financial markets. *Macroecon. Dyn.*, 4:170, 2000.
- [15] Thomas M Cover. Universal portfolios. *Mathematical finance*, 1(1):1–29, 1991.
- [16] David Morton de Lachapelle and Damien Challet. Turnover, account value and diversification of real traders: evidence of collective portfolio optimizing behavior. *New J. Phys*, 12:075039, 2010.
- [17] C. De Dominicis. Dynamics as a substitute for replicas in systems with quenched random impurities. *Phys. Rev. B*, 18:R3783, 1978.
- [18] Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.
- [19] J. D. Farmer. Market force, ecology and evolution. Technical Report 98-12-117, Santa Fe Institute, 1999.
- [20] Tobias Galla and Yi-Cheng Zhang. Minority games, evolving capitals and replicator dynamics. *Journal of Statistical Mechanics*, 2009:P11012, 2009.
- [21] Mauro Gallegati, Steve Keen, Thomas Lux, and Paul Ormerod. Worrying trends in econophysics. *Physica A: Statistical Mechanics and its Applications*, 370(1):1–6, 2006.
- [22] Francis Galton. Vox populi. *Nature*, 75:450–51, 1907.
- [23] Stanislao Gualdi, Giulio Cimini, Kevin Primicerio, Riccardo Di Clemente, and Damien Challet. Statistically similar portfolios and systemic risk. *arXiv preprint arXiv:1603.05914*, 2016.
- [24] W. Härdle and A. Kirman. Nonclassical demand: A model-free examination of price-quantity relations in the Marseille fish market. *Journal of Econometrics*, 67(1):227–257, 1995.
- [25] Giulia Iori. Avalanche dynamics and trading friction effects on stock market returns. *International Journal of Modern Physics C*, 10(06):1149–1162, 1999.
- [26] Trichet J.-C. Reflections on the nature of monetary policy non-standard measures and finance theory, November 2010. Speech by President of the ECB, Opening address at the ECB Central Banking Conference Frankfurt,.
- [27] Daniel Kahneman and Amos Tversky. Prospect theory: an analysis of decision under risk. *Econometrica*, 47:263, 1979.

- [28] Alan Kirman. Epidemics of opinion and speculative bubbles in financial markets. *Money and financial markets*, pages 354–368, 1991.
- [29] Alan Kirman. The economic crisis is a crisis for economic theory. *CESifo Economic Studies*, 56(4):498–535, 2010.
- [30] David Lando and Mads Stenbo Nielsen. Correlation in corporate defaults: Contagion or conditional independence? *Journal of Financial Intermediation*, 19(3):355–372, 2010.
- [31] J.M. Lasry and P.L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
- [32] Andrew Lo. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *J. Portfolio Management*, 30th Anniversary Issue, pages 15–29, 2004.
- [33] Terry Lohrenz, Kevin McCabe, Colin F Camerer, and P Read Montague. Neural signature of fictive learning signals in a sequential investment task. *Proceedings of the National Academy of Sciences of the United States of America*, 104(22):9493–9498, 2007.
- [34] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025, 2011.
- [35] Yannick Malevergne, Vladilen Pisarenko, and Didier Sornette. Testing the pareto against the lognormal distributions with the uniformly most powerful unbiased test applied to the distribution of cities. *Physical Review E*, 83(3):036111, 2011.
- [36] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin glass theory and beyond*. World Scientific, 1987.
- [37] Q. Michard and J.-P. Bouchaud. Theory of collective opinion shifts: from smooth trends to abrupt swings. *Eur. Phys. J. B*, 47:151–159, 2005.
- [38] P Read Montague, Brooks King-Casas, and Jonathan D Cohen. Imaging valuation models in human choice. *Annu. Rev. Neurosci.*, 29:417–448, 2006.
- [39] Vladimir Privman. *Nonequilibrium statistical mechanics in one dimension*. Cambridge University Press, 2005.
- [40] Thomas C Schelling. Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186, 1971.
- [41] Didier Sornette. Physics and financial economics (1776–2014): puzzles, ising and agent-based models. *Reports on Progress in Physics*, 77(6):062001, 2014.
- [42] James Surowiecki. *The Wisdom of Crowds*. Anchor Books, New York, NY, 2005.
- [43] Igor Swiecicki, Thierry Gobron, and Denis Ullmo. Schrödinger approach to mean field games. *Physical Review Letters*, 116(12):128701, 2016.
- [44] M. Tumminello, F. Lillo, J. Piilo, and R.N. Mantegna. Identification of clusters of investors from their real trading activity in a financial market. *New Journal of Physics*, 14:013041, 2012.

- [45] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [46] Vincent Weber and Florian Peres. Hedge fund replication: Putting the pieces together. *Available at SSRN 2202270*, 2013.
- [47] Jörgen W Weibull. *Evolutionary Game Theory*. MIT, Massachusetts, 1995.
- [48] Yi-Cheng Zhang. Why financial markets will remain marginally inefficient. In *Empirical Science of Financial Fluctuations*, pages 289–293. Springer, 2002.