

# Econophysics: a challenge to econometricians

Christopher A. Zapart<sup>a,b,1,\*</sup>

<sup>a</sup>*National Astronomical Observatory of Japan, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, JAPAN*

<sup>b</sup>*The Institute of Statistical Mathematics, 10-3 Midori-cho, Tachikawa, Tokyo 190-8562, JAPAN*

---

## Abstract

The study contrasts mainstream economics — operating on time scales of hours and days — with behavioural finance, econophysics and high-frequency trading, more applicable to short-term time scales of the order of minutes and seconds. We show how the central theoretical assumption underpinning prevailing economic theories is violated on small time scales. We also demonstrate how an alternative behavioural econophysics can model reactions of market participants to short-term movements in foreign exchange markets and, in a direct contradiction of the orthodox economics, design a rudimentary IsingFX automated trading system.

By replacing costly human forex dealers with banks of Field-Programmable Gate Array (FPGA) devices that implement in hardware high-frequency behavioural trading models of the type described here, brokerages and forex liquidity providers can expect to gain significant reductions in operating costs.

*Keywords:* econophysics, econometrics, Approximate Entropy, Ising spin model, high-frequency trading, foreign exchange currency markets

*PACS:* 89.65.Gh, 89.70.Cf, 75.10.Hk, 89.65.Gh, 89.75.-k

---

## 1. Introduction

The orthodox econometrics as well as the Random Walk Hypothesis (itself consistent with the Efficient Markets Hypothesis built upon Rational Expectations) treat logarithmic financial returns as a collection of i.i.d. (independently and identically distributed) random variables [1], which simplifies the use of statistical methods in finance. In essence, modern finance assumes that financial time series are random, investors make rational decisions and active short-term trading (as opposed to passive buy-and-hold investing) is referred to as futile

---

\*Corresponding author

*Email address:* `chris.zapart@nao.ac.jp` (Christopher A. Zapart)

<sup>1</sup>The research had been conducted whilst working at the Institute of Statistical Mathematics (ISM) in Tokyo, Japan. On the 1st of April 2014 the author moved from the ISM to National Astronomical Observatory of Japan. Consequently two addresses are listed.

"noise trading" [2]. In a perfect world this might well be true. However, back in the "real world" humans often act irrationally and, consequently, the i.i.d. assumption underpinning rational econometric models may not necessarily hold true. The irrationality of human behaviour is simply averaged out from mainstream economics and econometrics. Due to its inability to model erratic human behaviour, over the long run econometrics conveniently averages out irrationality from existence [3]. As once stated by the famous economist John Maynard Keynes, "in the long run we are all dead" [4], yet the readers of this article are most certainly alive. Similarly, extreme events and crashes still happen, and will keep coming with regularity for the foreseeable future since human nature never changes. Indeed, financial markets are driven by crowd behaviour, by fear and greed of irrational short-term traders, in other words by subjective psychological phenomena. Statistical testing whether or not financial markets are efficient is often obscured by making these tests conditional upon not necessarily correct parametric regression models, as has been made clear in [5, 6]. The inability of mainstream economists to anticipate short-term fluctuations of financial markets stands behind the recent rise in prominence [7] of the alternative behavioural finance, evidenced by the 2013 Nobel Memorial Prize in Economics co-awarded to the behavioural economist Robert J. Shiller. (In an interesting twist of fate, in 2013 the Nobel Memorial Prize was also awarded to Eugene F. Fama and Lars Peter Hansen, proponents of the mainstream econometrics which competes directly with behavioural finance. This only goes to show how difficult it is to form consensus on major issues in economics). Since psychology as well as behavioural finance both belong to the category of soft social sciences, they may not necessarily offer ready-to-use mathematical tools for building financial trading systems on top of them. Instead, their subjective findings need to be translated into objective trading rules, intended to execute automatically on computers without human interference. Statistical physics, and econophysics [8, 9, 10, 11] in particular may offer some models and tools for expressing in quantitative ways subjective human behaviour.

Human irrationality manifests itself in many ways. One effect acknowledged by the mainstream theoretical economists as well as practitioners is collective herding behaviour of traders. Using existing entropy analysis, in this article we reveal evidence of systematic violations of the i.i.d. assumption. It is often said that extraordinary claims require extraordinary evidence. Subsequently we also demonstrate how to build realistic foreign exchange trading systems based upon the idea of herding behaviour, as enforced through the use of the Ising spin model, common in statistical physics, and adapted to the financial domain by the author. The idea of applying the Ising model to financial markets or social phenomena is not new [12, 13]. The large body of existing econophysics literature, reviewed in for example [12], tends to focus on running artificial agent-based simulations with realistic supply/demand-based artificial price formation mechanisms, in order to reproduce so-called "stylised facts" (volatility clustering, fat tails etc.). In contrast, the IsingFX model described in this paper does not include any price formation mechanism. Nor is it used to generate any artificial price time series, to be compared against the dynamics of a real

market. Instead the spins (artificial traders) within the IsingFX model react to real forex prices streamed in real-time using a C/C++ FIX API connection to the author's foreign exchange trading account. As the output of the real-prices-driven IsingFX, the net lattice magnetisation is translated into BUY/SELL trading decisions, ready to be transmitted to the forex market using the FIX protocol (an industry standard).

## 2. Entropy analysis

Approximate Entropy (ApEn), being “a model independent measure of sequential irregularity” [14], has been employed in this study to demonstrate beyond reasonable doubt that high frequency foreign exchange time series do exhibit certain sequential regularities incompatible with the i.i.d. assumption made by mainstream econometricians [1]. A recent application of ApEn to study speculative bubbles conditions in Tunisian and French stock markets can be found in [15]. Itself non-parametric, Approximate Entropy is also capable [16] of either endorsing or rejecting parametric econometric models such as ARIMA or GARCH [1].

Tick data collected from the foreign exchange market during the three-week period between 10th and 29th March 2014 are used to construct time series of logarithmic returns  $\log s_t - \log s_{t-1}$  for ten selected currency pairs, where  $s_t$  are middle prices between bid and ask quotes. Employing a sliding window of the length  $N = 129$ , let us assume a logarithmic returns sequence  $\{x_i\}$ ,  $i = 1 \dots N$ . The value of Approximate Entropy  $ApEn(m, r, N)$  for a particular sequence (sliding window) is calculated using the following parameters:  $m = 2$  and  $r = 0.4 \times mad$ , where  $mad$  denotes a mean absolute deviation of  $\{x_{i+1} - x_i\}$  (robustness to outliers), as opposed to the standard deviation used in the original Approximate Entropy measure [14, 16]. Based on the sequence  $\{x_i\}$ , delay vectors of the length  $m$  are constructed: the  $i$ th delay vector  $\mathbf{x}(i) = [x_i, x_{i+1}, \dots, x_{i+m-1}]$  and the  $j$ th delay vector  $\mathbf{x}(j) = [x_j, x_{j+1}, \dots, x_{j+m-1}]$ . Let us define a quantity  $C_i^m(r)$  to be

$$C_i^m(r) = \left( \text{the number of } \mathbf{x}(j) \text{ such that distance}(\mathbf{x}(i), \mathbf{x}(j)) < r \right) / (N - m + 1)$$

with the distance measure such that  $\text{distance}(\mathbf{a}, \mathbf{b}) = \max_{k=1 \dots m} |\mathbf{a}_k - \mathbf{b}_k|$  and  $\text{distance}(\mathbf{x}(i), \mathbf{x}(i)) = 0$ . Then Approximate Entropy for the sequence  $\{x_i\}$  is defined to be

$$ApEn(m, r, N) = \Phi^m(r) - \Phi^{m+1}(r)$$

where  $\Phi^m(r) = (N - m + 1)^{-1} \sum_{i=1}^{N-m+1} \log C_i^m(r)$ .

In the world of algorithmic high frequency trading three weeks is more than enough to “make or break” an algorithm. Each logarithmic returns time series contains over one million price ticks coming at time intervals ranging from sub-second to few seconds, depending on the time of the day. Figure 1 shows an experimental implementation of the algorithm to compute the Approximate Entropy using Xilinx Field-Programmable Gate Array (FPGA) technology. The

real reason for making an extra effort to implement ApEn in hardware was to practise bit-level digital hardware design in VHDL (VHSIC Hardware Description Language). In recent years the FPGA technology has found widespread use within the high-frequency trading [17] community as well as global investment banks (for example JP Morgan [18, 19]), desperate to win the speed race to reduce trading latencies.

In order to ascertain the presence of sequential regularities, the logarithmic returns series is randomly reshuffled — destroying any sequential information — and Approximate Entropy is then re-calculated for the randomised time series. Had the original time series been i.i.d. (identically and independently distributed), there would be little difference in Approximate Entropies (measuring sequential irregularities) obtained for the original and randomly reshuffled time series. Figure 2 shows histograms of Approximate Entropies obtained from the original and reshuffled logarithmic returns time series. For the ten currency pairs under consideration, in all but one case original time series clearly exhibit lower entropy levels than what would be expected by chance, indicative of the presence of sequential regularities in high frequency financial data. One notable exception stands out: EUR/CHF. One could speculate that the on-going exchange rate floor maintained by the Swiss National Bank distorts free markets and alters natural speculative positioning in the EUR/CHF currency pair.


When the data sampling frequency is lowered (for example by taking every second, fifth, tenth etc. datum from the tick time series), differences between original and reshuffled entropy histograms gradually diminish. Figure 3 plots  $\chi^2$ -distances between histograms as a function of a sampling scale (distances are calculated using a computer science method described in [20]). Distances between histograms from the Figure 2 show a declining tendency as sampling intervals increase. The degree to which financial data satisfies the i.i.d. assumption seems to depend on the time scale at which the price is sampled. At small time scales, corresponding to seconds and minutes, the foreign exchange data exhibits more regularity than implied by pure chance. Private retail traders are unlikely to be able to exploit sub-second market inefficiencies due to prohibitive trading costs (large bid/ask spreads). However, high frequency trading firms and market makers (liquidity provider banks) operating at sub-second time scales, at which currency markets are not random, are subject to substantially lower trading costs, enabling them to make profits at the expense of longer-term investors and small market players.

After repeatedly reshuffling the original data using different pseudo-random number sequences, the histogram differences shown in Figures 2 and 3 have been found to be insensitive to the initial choice of pseudo-random number generator (PRNG) seeds. The large number of samples combined with a relatively high quality of the Mersenne Twister 19937 generator used in this study provides a plausible explanation for such a finding. The value of ApEn depends on taking logarithms of average numbers of delay vectors that appear similar. Random reshuffling using a high-quality PRNG does not seem to alter to a large enough extent the average counts of similar vectors, especially as the number of counts is large. In theory there is nothing to prevent random reshuffling from producing a

inclusive of PCI Express data  
transfers, FPGA is 15 times faster

Kintex-7 XC7K325T FPGA  
100 MHz, 3.5 Watts

Intel Core i7 3770S CPU  
3.1 GHz, 65 Watts



VS.




Register Transfer Level  
VHDL hardware design

Intel C/C++ compiler  
4/8HT cores, OpenMP  
data fits within 8MB cache

computation speed: 15 times faster  
performance per watt: 280 times better

Wolfram Mathematica



MathLink C/C++ interface  
PCI Express bus

Xilinx FPGA




Figure 1: **FPGA hardware acceleration** Computation of the Approximate Entropy for  $N = 129$  and  $m = 2$  has been offloaded onto a custom digital circuit, designed by the author in VHDL (VHSIC Hardware Description Language) and implemented in a Xilinx Kintex-7 FPGA (Field-Programmable Gate Array) device. The hardware-accelerated FPGA implementation has been found to be over 15 times faster compared with a C/C++ code running on an Intel CPU whilst consuming an order of magnitude less power. As an invaluable exercise in CPU-FPGA integration, the author has also created a C/C++ MathLink/PCI Express interface between Wolfram Mathematica and Xilinx FPGA, enabling one to stream financial time series from Mathematica directly to FPGA and have the FPGA send back calculated entropy values.

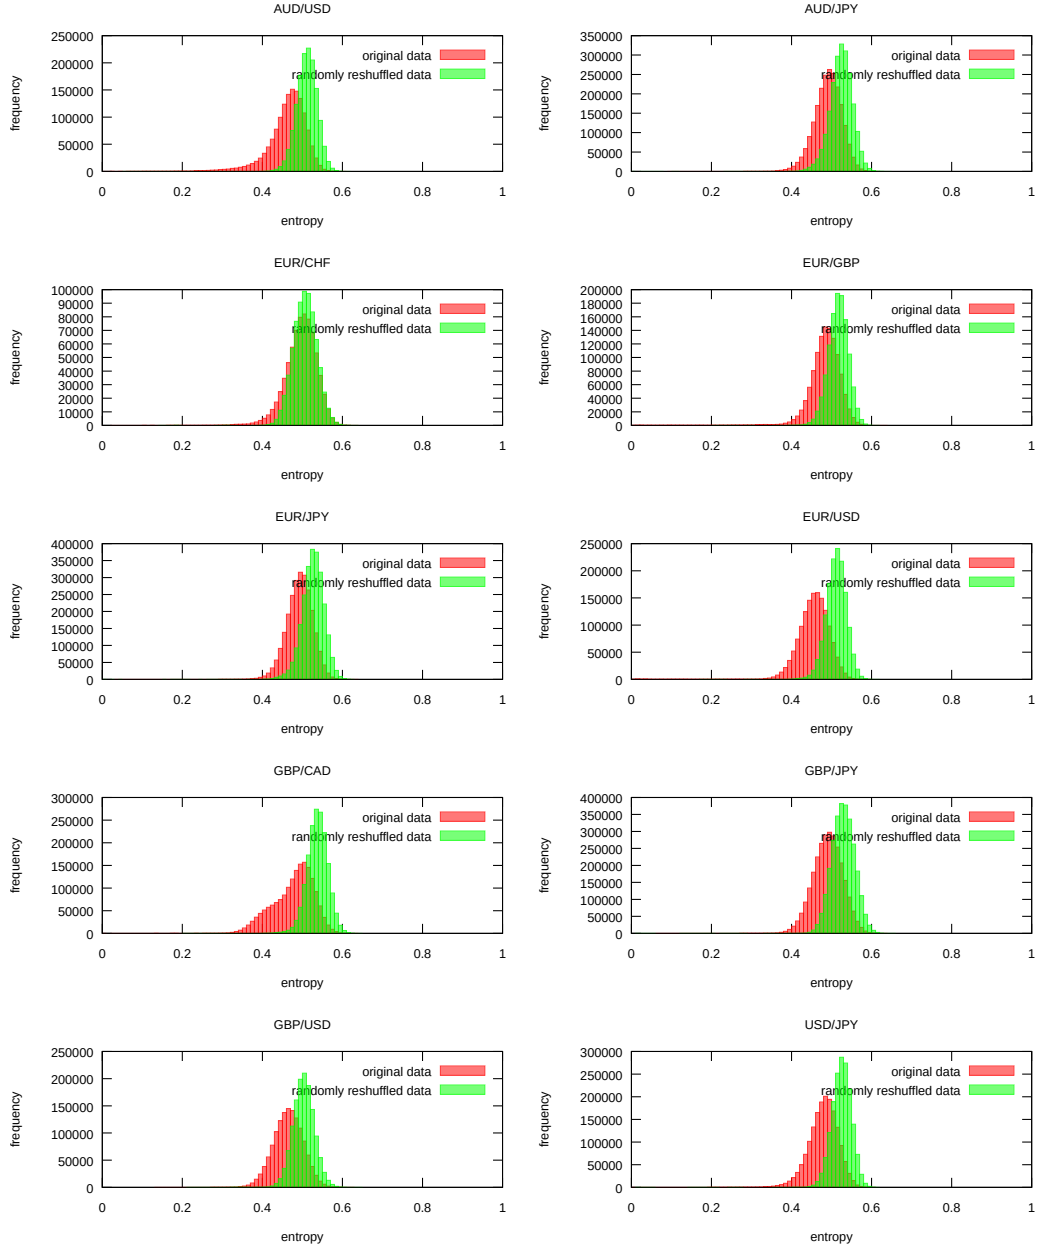


Figure 2: **Entropy analysis** Histograms of Approximate Entropy obtained for ten selected currency pairs using tick-level time series (sub-second time resolution). Lower entropy values are indicative of increased sequential regularities present in the time series. A notable exception is the EUR/CHF currency pair.




Figure 3: **Histogram differences**  $\chi^2$ -distances between histograms of Approximate Entropy seen in Figure 2 calculated using an empirical histogram comparison method described in [20].  $scale = 1$  corresponds to sampling data at each tick,  $scale = 10$  implies using every 10th sample, and so on. The distance calculation method bears no relation to statistical  $p$ -value  $\chi^2$  tests. In addition, histogram differences have been found to be insensitive to reshuffling original data using different pseudo-random number generator seeds, which makes it difficult to plot any error bars (different experimental runs result in the same histogram differences for a given currency pair). Instead a natural variability amongst ten different currency pairs provides a proxy to error bars.

new time series that is very similar to the original financial time series. However, in practice one would need to wait prohibitively long before finding significant deviations from average similarity counts.

### 3. IsingFX

Is it possible to exploit market inefficiencies and sequential regularities revealed by the entropy analysis? The following section attempts to answer this question. Short-term fluctuations in foreign currency prices arise mainly as a result of three factors: a random arrival of news, scheduled releases of economic indicators that influence interest rate differentials, and internal dynamics of the market (interactions between market participants) caused by for example stop-loss hunting, attempts to trigger binary options, positioning by big players in the direction of a prevailing trend, short-term mean-reversion etc.. Whilst forecasting the impact of news is rather difficult, after the news has been released it may be feasible to model human reactions to price changes and/or incurred profits/losses resulting from a random arrival of news, as illustrated in Figure 4. To emphasise, the author does not claim to be able to predict the arrival of shocks and dislocations in financial assets. Instead behavioural econophysics attempts to model traders' reactions to profits and losses incurred due to shocks. After a release of economic indicators, those traders caught on the wrong side are often forced to change the direction of their positions due to the use of leverage that grossly amplifies losses, breaching their risk limits.

Econometricians tend to ignore the microscopic interactions between market participants, preferring to focus on the “big picture” (macroeconomics, top-down stochastic processes), although there are notable exceptions [21]. In contrast, behavioural econophysics borrows from the Ising spin model (common in statistical physics) to model a subset of human interactions. It does so by assuming a 2D square lattice, an example of which appears in Figure 5, containing  $N$  imaginary traders that form an artificial foreign exchange market. Each trader makes either *buy* or *sell* decisions subject to minimising the following energy (cost) function:

$$h_i = -\alpha \underbrace{\sum_{j \in N_i} J_{ij} S_i S_j p_j}_{\text{copy successful traders}} - \underbrace{\beta \min(p_i, 0)}_{\text{cut own losses}} + \underbrace{\gamma \max(p_i, 0)}_{\text{take profits}}, \quad \alpha, \beta, \gamma \in \mathbb{R}^+,$$

where  $h_i$ ,  $i = 1 \dots N$  is the  $i$ th trader's energy function,  $\alpha, \beta$  and  $\gamma$  are tunable model parameters,  $J_{i,j} \in [-1, 1]$  controls the strength of interactions between the  $i$ th and  $j$ th traders,  $N_i$  denotes an immediate neighbourhood of the  $i$ th trader (readers are assumed to be familiar with the Ising spin model [22]),  $S_i, S_j \in \{-1, 1\}$  encode binary *sell* or *buy* decisions, respectively, and  $p_i$ ,  $i = 1 \dots N$  holds current profits/losses. By changing the sign of  $J_{i,j}$  one can simulate either herding (trend following) or minority game-like mean reversion. The model employs a Monte Carlo method in order to minimise the total energy of the 2D lattice. Upon receiving a new currency quote from the forex broker, one



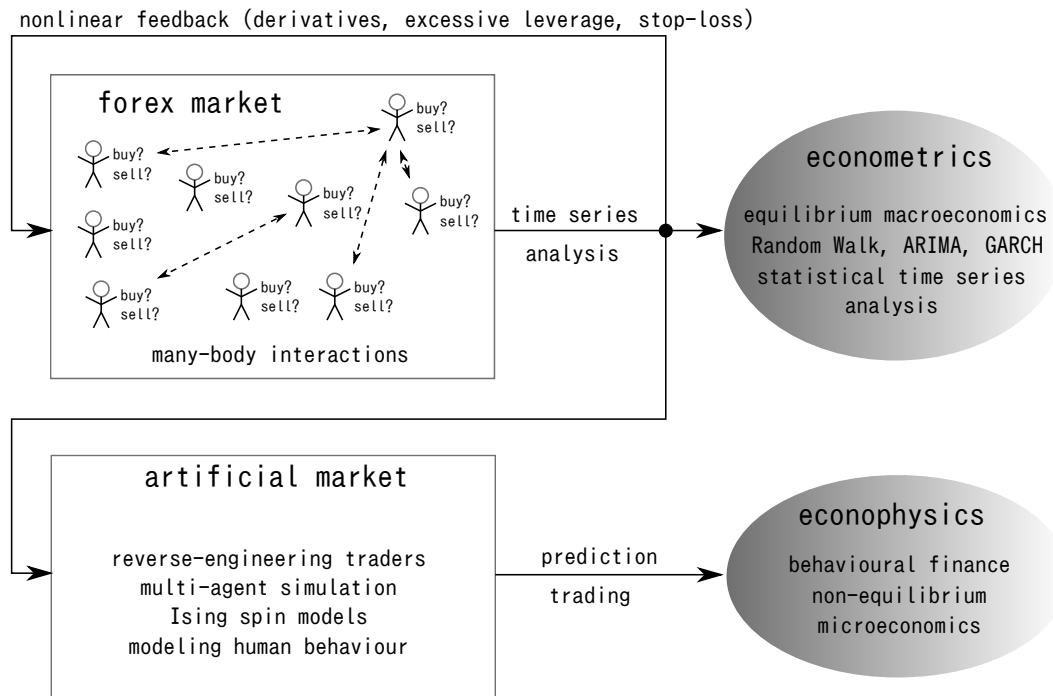



Figure 4: **Microscopic interactions** Foreign exchange price changes caused by interactions between market participants. Econometrics observes the resulting time series (*effects*) whilst ignoring the underlying *causes* of those price changes. Using a medical analogy, this is equivalent to treating symptoms instead of going after root causes of a disease.

2011/05/27 16:59:39 USDJPY 80.79 UP 363 DOWN 537 NET -174



793689 E 64.49 kBT 0.10 CUMULATIVE -5.859 dE>0 FLIPPED 45%

Figure 5: **2D Ising lattice** A representative visualisation of a 2D Ising lattice adapted to forex trading. Each square represents an imaginary trader, with buy and sell decisions denoted by different colours.

trader randomly selected from the 2D lattice undergoes a trial spin flip in an attempt to minimise its energy  $h_i$ . If, upon flipping a spin (changing the position from for example *buy* to *sell*), the energy  $h_i$  decreases, the spin flip is accepted. On the other hand, if the energy change  $dE$  is greater than zero, a spin flip is accepted with a probability given by the Boltzmann factor  $\exp(-dE/k_B T)$ , where  $k_B$  denotes the Boltzmann constant and  $T$  is the temperature (akin to temperature used in Simulated Annealing). However, in a departure from the standard Simulated Annealing Monte Carlo method, the algorithm automatically *increases* or *decreases* the temperature  $T$  so as to maintain the acceptance probability of  $dE > 0$  spin flips constant on average. The author has chosen the name “IsingFX” for this model.

Econometricians usually pay a great deal of attention to model parameter estimation. Unfortunately no amount of fitting distributions to past data prior to the stock market meltdown in 2008 would have revealed the severity of the incoming downturn. In contrast, the model described in this article does not

involve any sort of fitting statistical models to past time series nor optimising Sharpe Ratio-based fitness functions. Hence it completely avoids the problem of over-fitting that plagues econometrics. Once the  $\alpha, \beta$  and  $\gamma$  parameters governing the desired behaviour of artificial traders have been decided upon, the IsingFX model does not require any further tuning nor fitting to past data.

Simulated trading performance of the IsingFX model operating in a tick-level high frequency mode is shown in the Figure 6. Cumulative profits/losses are averaged over all  $32 \times 32 = 1024$  traders. Positive net magnetisation of the 2D spin lattice corresponds to the majority of traders holding a *long* (buy) position. Conversely negative magnetisation implies a net *short* (sell) position. Imaginary traders change the trading direction at a very high frequency, which would be impossible to match by average retail traders paying high bid/ask transaction costs. The value of  $k_B T$  is automatically adjusted in order to keep the  $dE > 0$  spin flip acceptance probability constant around 0.2. To do so we keep track of log-odds [23] of spin flips given  $dE > 0$  using a statistical ensemble [22] consisting of 128 independent IsingFX models. The author does not recommend using time averages for estimating the  $dE > 0$  spin flip conditional probability. Given the non-stationary nature of financial markets, the correct approach [24, 25] is to use statistical ensemble averages, as illustrated in Figure 7. Time averages are only applicable to cases in which the stationarity assumption holds [24, 25]. The 0.2 conditional probability of  $dE > 0$  spin flips corresponds to the log-odd value of  $-1.39$ , which the control algorithm can be seen to maintain on average, as seen in the Figure 8. Due to the stochastic nature of financial markets the log-odds are seen to oscillate around the target value of  $-1.39$  as  $k_B T$  is being continuously adjusted. There is no theoretical basis for targeting specifically the value of 0.2 for the  $dE > 0$  spin flip acceptance probability. Being a probability, its value must lie between 0 and 1. Setting it to zero would prevent Simulated Annealing-style hill-climbing (exploration) of the energy landscape. However, setting it to 1 would be equally counter-productive as the behaviour of the IsingFX model would have become purely random. Therefore a trading system designer has to make a heuristic choice of a value between 0 and 1. The entity  $k_B T$  can probably be interpreted as a proxy for *intrinsic volatility*, quickly rising during highly volatile times and falling promptly to a small non-zero baseline level after the shocks have passed.

As an alternative to adjusting  $k_B T$  in accordance with the probability of  $dE > 0$  flips, statistical physics also offers a temperature-less *Demon Monte Carlo* spin flip dynamics. Compared to the variable  $k_B T$  scheme, an experimental Demon Monte Carlo version of IsingFX (appropriately named DemonFX) has been found to offer similar trading performance. However, the VHDL FPGA hardware implementation [26] of DemonFX is easier to realise on a practical basis than the IsingFX model discussed in this paper. The Figure 9 illustrates a VHDL hardware design of a 32-core high-frequency trading chip implemented using a Kintex-7 Xilinx FPGA.

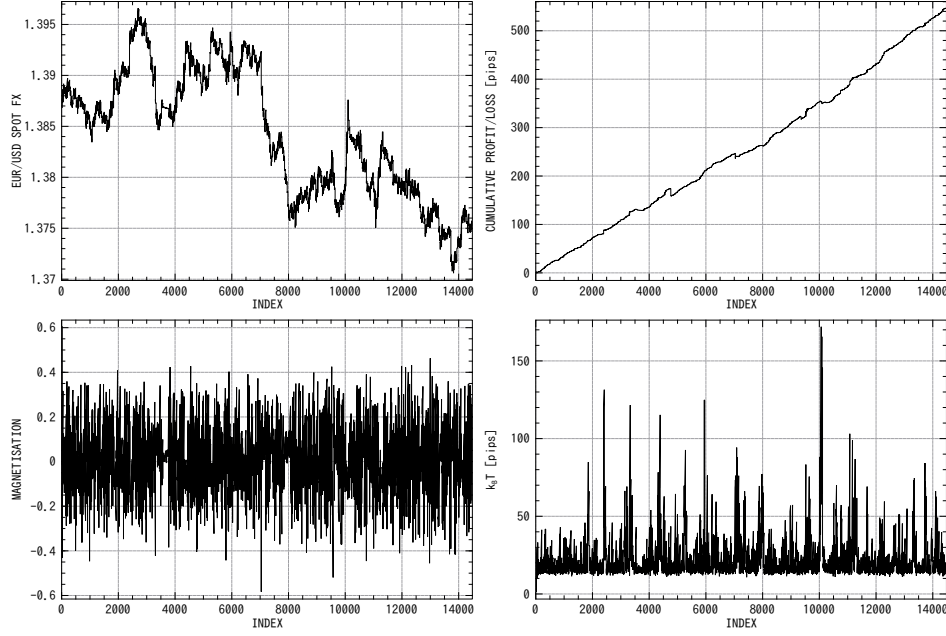


Figure 6: **Simulated trading** EUR/USD high frequency trading performance of the IsingFX model plotted every 100th tick for the period between 10th and 29th March 2014. Only high frequency trading firms — able to trade within the bid/ask spread — are capable of approaching trading profits seen in the “cumulative profit/loss” curve. Private retail traders, subject to large transaction costs, are effectively prevented from being able to rebalance their positions with the frequency seen in the “magnetisation” curve. Position sizing is proportional to the net magnetisation of the 2D Ising lattice, with the trading direction (*long* or *short*) given by the sign of the magnetisation.




Figure 7: **Statistical ensemble** Mainstream econometrics (Kalman filtering, moving averages, historical volatility) implicitly relies on time averages to estimate the mean value of  $x$  at time  $t$  using past values  $x(t), x(t-1), x(t-2), \dots$ , which introduces undesirable time lags. In contrast, econophysics utilises the *statistical ensemble* approach in which an ensemble of  $N$  independent models (or re-runs of an experiment) provides an instantaneous (intrinsic) mean value of  $x(t)$ .

#### 4. Conclusions

As an immediate practical application, the IsingFX can contribute towards reducing operating costs and increasing profits at brokerages and forex liquidity providers. Foreign exchange brokers typically internalise customer order flow (match buys with sells). The remaining imbalance (Net Open Position or NOP) would normally be hedged in the interbank forex market at some future point. The exact timing of hedging the NOP is left to the discretion of skilled forex dealers. Some brokerages may be trading against their customers by choosing not to offset the NOP in the hope that their customers would be forced to close their positions at a loss later on, thus helping to reduce the NOP without resorting to the interbank market. The IsingFX model can help dealers decide on the optimum timing of hedging the NOP. In an extreme case, expensive human forex dealers could conceivably be replaced by banks of Field-Programmable Gate Array (FPGA) devices implementing in hardware high-frequency behavioural trading models of the type described here. A prototype VHDL FPGA implementation of IsingFX/DemonFX has been discussed in [26].

The IsingFX model conveniently bypasses the problem of over-fitting to past data by replacing fitting econometric time series models with modelling the underlying behaviour of human traders. However, an initial guidance from an experienced forex trader and/or trading system designer is still required in order to decide which aspects of human behaviour should be expressed using Ising spins.

When asked a question “Is the Random Walk Theory correct?”, it seems the answer would depend on what time scale one looks at. At medium to large time scales the Random Walk probably provides a reasonable first approximation to movements of financial markets. However, on the scale of minutes




Figure 8: **Constant log-odds** A zoom-in look at the results from Figure 6. To estimate the probability of  $dE > 0$  spin flips the model keeps track of log-odds of  $dE > 0$  flip events measured using a statistical ensemble consisting of 128 independent IsingFX models.




Figure 9: **DemonFX ensemble in hardware** A Register Transfer Level VHDL design of a parallel 32-core DemonFX statistical ensemble. The floating-point implementation of 32 DemonFX cores fits comfortably within the Kintex-7 XC7K325T Xilinx FPGA. By replacing floating-point with a fixed-point arithmetic it is quite possible to fit more parallel cores within the same FPGA device. Alternatively a newer, much bigger Kintex UltraScale-class FPGA can accommodate 64 or 128 floating-point DemonFX cores.

and seconds — the domain reserved mainly to high-frequency trading firms and market making liquidity providers — the markets seem far from being efficient and random; consequently the Random Walk Theory cannot be said to explain how foreign exchange markets work. In contrast with econometrics, which makes randomness and unpredictability of financial markets as its central tenets, non-parametric behavioural econophysics of the kind presented here is capable of making short-term directional calls that can be exploited to design high frequency trading systems. The author feels the time is right to recognise econophysics as a refreshing alternative to overly dogmatic econometrics.

## References

- [1] R. S. Tsay, Analysis of Financial Time Series, 3rd Edition, WILEY SERIES in PROBABILITY and STATISTICS, John Wiley & Sons, New Jersey, 2010.
- [2] F. Black, Noise, The Journal of Finance 41 (3) (1986) 529–543.
- [3] R. C. Merton, Optimum consumption and portfolio rules in a continuous-time model, Journal of Economic Theory 3 (1971) 373–413.
- [4] J. M. Keynes, A Tract on Monetary Reform, 1923, Ch. 3.
- [5] L. H. Summers, Do we really know that financial markets are efficient?, Working Paper 994, National Bureau of Economic Research (September 1982). doi:10.3386/w0994.  
URL <http://www.nber.org/papers/w0994>
- [6] L. H. Summers, Does the stock market rationally reflect fundamental values?, The Journal of Finance 41 (3) (1986) 591–601.
- [7] M. J. Casey, FX HORIZONS: Move over economists, time to give physicists a turn, The Wall Street Journal (2013) Markets & Finance Blog.
- [8] R. N. Mantegna, H. E. Stanley, AN INTRODUCTION TO ECONOPHYSICS Correlations and Complexity in Finance, Cambridge University Press, 2000.
- [9] H. Takayasu, The Application of Econophysics, Springer Japan, 2004, Proceedings of the Second Nikkei Econophysics Symposium.
- [10] H. Takayasu, Empirical Science of Financial Fluctuations: The Advent of Econophysics, Springer, 2001.
- [11] M. Buchanan, What has econophysics ever done for us?, Nature Physics 9 (317).
- [12] D. Sornette, Physics and financial economics (17762014): puzzles, Ising and agent-based models, Reports on Progress in Physics 77 (6) (2014) 062001.  
URL <http://stacks.iop.org/0034-4885/77/i=6/a=062001>



- [13] C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics, *Rev. Mod. Phys.* 81 (2009) 591–646. doi:10.1103/RevModPhys.81.591.  
URL <http://link.aps.org/doi/10.1103/RevModPhys.81.591>
- [14] S. M. Pincus, Approximate entropy as a measure of system complexity, *Proceedings of the National Academy of Sciences USA* 88 (6) (1991) 2297–2301.
- [15] I. Mahmoud, K. Naoui, H. Jemmali, Study of speculative bubbles: The contribution of approximate entropy, *International Journal of Economics and Financial Issues* 3 (3) (2013) 683–693.
- [16] S. M. Pincus, R. E. Kalman, Irregularity, volatility, risk and financial market time series, *Proceedings of the National Academy of Sciences USA* 101 (38) (2004) 13709–13714.
- [17] C. Leber, B. Geib, H. Litz, High frequency trading acceleration using FPGAs, in: *Field Programmable Logic and Applications (FPL)*, 2011 International Conference on, IEEE, 2011, pp. 317–322.
- [18] O. Mencer, S. Weston, Computational acceleration of credit and interest rate derivatives, *Global Derivatives Trading and Risk Management*.
- [19] W. Wans, Black-Scholes in hardware, *Overload* 8.
- [20] O. Pele, M. Werman, The quadratic-chi histogram distance family, in: K. Daniilidis, P. Maragos, N. Paragios (Eds.), *Computer Vision ECCV 2010*, Vol. 6312 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2010, pp. 749–762. doi:10.1007/978-3-642-15552-9\_54.  
URL [http://dx.doi.org/10.1007/978-3-642-15552-9\\_54](http://dx.doi.org/10.1007/978-3-642-15552-9_54)
- [21] Y. Hashimoto, T. Ito, T. Ohnishi, M. Takayasu, H. Takayasu, T. Watanabe, Random walk or a run: Market microstructure analysis of the foreign exchange rate movements based on conditional probability, *Working Paper 14160*, National Bureau of Economic Research (July 2008).  
URL <http://www.nber.org/papers/w14160>
- [22] F. Reif, *Fundamentals of Statistical and Thermal Physics*, McGraw-Hill, 1965.
- [23] B. Nessler, M. Pfeiffer, W. Maass, Reward-modulated Hebbian learning of Bayes optimal decisions, *Neural Computation* 22 (6) (2010) 1399–1444.
- [24] J. L. McCauley, Thermodynamic analogies in economics and finance: instability of markets, *Physica A* 329 (2003) 199–212.
- [25] J. L. McCauley, *Stochastic Calculus and Differential Equations for Physics and Finance*, Cambridge University Press, 2013.

- [26] C. A. Zapart, Reconfigurable computing: towards an FPGA hardware implementation of multi-agent simulations (translated from Japanese), Research Collaboration Report 292, The Institute of Statistical Mathematics (March 2013).  
URL [http://kyodo.ism.ac.jp/search/php/r\\_kyodo2.php?no=292](http://kyodo.ism.ac.jp/search/php/r_kyodo2.php?no=292)