Eur. Phys. J. Special Topics **225**, 3243–3254 (2016)
© EDP Sciences, Springer-Verlag 2016
DOI: 10.1140/epjst/e2016-60123-4

THE EUROPEAN PHYSICAL JOURNAL SPECIAL TOPICS

Regular Article

Financial price dynamics and agent-based models as inspired by Benoit Mandelbrot

Blake LeBaron[a]

International Business School, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453–2728, USA

Received 18 April 2016 / Received in final form 15 September 2016
Published online 22 December 2016

Abstract. This short note draws some connections between Mandelbrot's empirical legacy, and the interdisciplinary work that followed in finance. Much of this work is now labeled econophysics, but some has always been more in the realm of economics than physics. In a few areas the overlap is even becoming quite complete as in market microstructure. I will also give some ideas about the various successes and failures in this area, and some directions for the future of agent-based modeling in particular.

1 Introduction

Agent-based financial and economic modeling has been greatly influenced by the work of Benoit Mandelbrot. Now over 50 years old his work still influences researchers from many different backgrounds who have become interested with some of the most basic problems in finance.

In this short note I will look at the current state of agent-based modeling through some of its ties to Mandelbrot's work, and other more recent results. As with Mandelbrot himself, much of this work has occurred in an interdisciplinary realm between finance, physics, mathematics, and computer science (with a little biology too). It would be foolish to attempt a survey of this literature, and would go well beyond the bounds of this note. Therefore, my goal is more to poke the reader into new directions and literature that is just appearing, than to summarize the past. I will also present some thoughts on best practice for building financial models, and some general directions.

Since this is not a survey of econophysics it is important to point out some recent excellent surveys which are available. These include [18,22], and [86], which come from a physics background. Many of the most important empirical contributions in this area have come in the subfield of finance known as market microstructure. This area looks at the dynamics of financial markets at detailed trade by trade level. Again very nice surveys of this area are [19,31]. Also, I will be presenting a lot of commentary on agent-based models which appear more in the economics world, but have an interdisciplinary overlap. The reader may be interested in three surveys from the same volume [26,60,77]. All bring a slightly different perspective to the field.

[a] e-mail: blebaron@brandeis.edu

Therefore, my goals in this note are to acknowledge the legacy of Mandelbrot's work. I will then discuss some of the empirical impact of interdisciplinary work that has come after him. Here, I will comment on what has been most influential in our understanding of market dynamics, and where ideas from the interdisciplinary realm have been accepted into mainstream financial thinking. Finally, I will be discussing agent-based models in general, both from a standpoint of what they have accomplished, but more important, I will try to lay out a prediction of where they may be going.

2 Mandlebrot's legacy and empirical econophysics

2.1 Return distributions

Mandelbrot changed how we view financial time series in a deep way. His early work discovered several important features which are still not completely understood today. For many of us it is probably true that we touch on his work, at least tangentially, every day.

I need to first do one quick definition for a possible interdisciplinary readership. In the world of finance we are most often more interested in returns, and not prices per se. If an asset price at time t is given by P_t, then its one period return would be

$$R_{t+1} = \frac{P_{t+1} - P_t}{P_t}$$

or in logs where $p_t = \log(P_t)$,

$$r_{t+1} = p_{t+1} - p_t = \log(P_{t+1}) - \log(P_t).$$

For relatively short time horizons and small returns one gets $r_t \approx R_t$, so for most of the commentary in this note the units are irrelevant.

Mandelbrot's first critical result showed that the returns of various assets at relatively high frequency (daily) were not Gaussian, but were fat tailed. We simply take this fact for granted now, but in 1963 [79] was revolutionary to the emerging field of quantitative finance.[1] The importance of this result was somewhat ignored by the finance community. Financial economists and practitioners were more interested in approximate pictures of the entire distribution, and not the tails. They often felt a Gaussian approximation was not bad, and also relied on a key result that longer horizon returns do get closer to a normal distribution. The feeling was that the average investor could ignore these disturbing features. The appearance of options markets, and the advent of modern risk management changed this. In both cases the entire distribution comes into play, and in the latter case the distribution of extreme events is critical to determining behavior in stressed situations. By the 1990's Gaussian approximations were showing their many weaknesses, and empirical and theoretical work on tail behavior began to take off.

Much empirical work in econophysics has centered around better understanding of these facts. It is often gone places that financial economists were unwilling to go, or simply not looking at. It also operates outside the bounds of certain constraints in the more traditional economics world which is often skeptical of broad atheoretic empirical work. A research paper in economics almost always needs to start with a relatively well understood model, and then moves from there to empirical work which is connected to the model. At times this approach can be useful, but it often leads

[1] This result is also in [45] as well. To get a general feeling of where researchers were at that time see [33].

to a kind of local search for empirical facts. The empirical side of econophysics is not subject to these constraints.

A critical starting point for empirical econophysics was the publication of [80]. This paper is an example of several of the important extensions to Mandelbrot's work which are common to other papers in this field. First, it moved to much higher frequency data since that had started to become available at the time. Also, it takes the quantitative estimation of power laws much more seriously than Mandelbrot did in his earlier work. The actual value of a tail exponent, measuring the tail shape was now something of great interest. It appeared we could estimate it very reliably, and it also appeared to be relatively stable across many return frequencies from 5 minutes through nearly two weeks.

The usefulness in power-law behavior in returns has certainly been under appreciated in the economics and finance worlds. Given that a distribution for returns is not normal, power-law exponents (or more formally tail exponents) give us an estimate of probabilities of extreme return events. From a risk management perspective these can be extremely important. Also, a related result is that higher order moments depend on the tail exponent. For example a tail exponent of 3 implies that all moments $E(X^k)$ for $k \geq 3$ do not exist. Given that exponents in the range of 3–4 have been commonly estimated in the literature, we should be wary of many common measures such as kurtosis, and some forms of volatility estimates, all of which require the existence of moments beyond 3. Though traditional asset pricing models have generally ignored tail exponents, they are beginning to acknowledge their usefulness as in [62].

What are the weaknesses present in this vast literature which sits mostly in the physics world? One is that it has generally failed to connect itself to the related area of extreme value theory which looks at the behavior of distribution tails [40]. A stronger unification with this area would have been useful in the early parts of this line of research. Another would be a greater appreciation that the tail exponent itself is a complicated object to estimate. Many papers in this area probably over estimate the precision of their estimates. In practical risk management applications one can often be frustrated with the abilities to estimate tail behaviors when only limited amounts of data are available. Estimating tail exponents is a classic bias/variance trade off problem. Moving too far out into the tail reduces the amount of data available (variance increases), but not moving far enough moves away from the log/log scaling region (bias increases). Balancing these can be tricky.[2]

Another interesting question is exactly what are these distribution tails about? Mandelbrot argued that they are part of the unified distribution generated by the price dynamic. There is no different process going on in the tail. I think this clean notion would be consistent with the spirit of much of econophysics research, and also all agent-based modeling which is thinking about a unified model for heterogeneous agent behavior and price movements. In finance an alternative idea is represented by the idea of jump processes which model prices as a diffusion with occasional large jumps drawn from another distribution [82]. This idea has obviously been around for a long time, but it is making a kind of empirical comeback in finance. First, high frequency modeling has given us new econometric technology to separate the jumps from the diffusion components.[3] If jumps are somehow special then days with jumps should be different. There is some evidence that this is true both for returns, and the estimation of correlations across returns. In [42] the author shows that leaving out a small number of large move days greatly reduces the performance of long term

[2] [29] is a nice summary of power-law estimation technologies. [35] presents evidence on how difficult tail estimation can be in real-world risk management situations.

[3] This comes from the field of realized volatility. Recent work on the econometrics of jumps is in [8].

investment strategies. Further evidence on this can be found in [84] who show that there are almost two different types of markets, one when announcements occur and the other without. Announcement days yield larger returns, and more informative co-movements across assets. It is also possible that most large moves can be connected to more endogenous changes at the level of detailed trading dynamics [48]. Although there has been a lot of progress, the question of where large price moves come from is still not answered.

2.2 Volatility persistence

Not as many people realize Mandelbrot's second contribution. In his early papers he noted in passing that beyond the fat tailed distributions there was a second unusual piece of behavior in the returns series he was looking at. Large returns (of either sign) tended to follow other large returns. In other words, the magnitudes of returns are persistent. This notion led to discoveries of markets moving through periods of high and low volatility (or turbulence). It is interesting that this important feature did not get fully recognized until much later.[4] The reason for the slow digestion of this fact is probably similar to previous one. It needed to wait for a quantitative finance community that was highly motivated to measure and understand the dynamics of changing conditional distributions.[5]

The econophysics literature has extended our knowledge of this feature as well. It stressed the stylized power law decays in autocorrelations of the absolute values of returns, $|r_t|$, and showed these to be robust across many asset classes and length scales.[6] The fact that volatility dynamics appear to generate a relatively regular dependence decay that is strongly supportive of a long memory process is a major fact. I think the impact of this in our thinking about learning models is still under appreciated in the world of building agent-based models.

Again, the beauty of this mathematical result may also be at odds with some empirical facts. The calendar often appears as a driver of volatility itself as there are strong seasonalities over the day and at longer horizons. One of the most subtle of these involves asymmetries across time horizons. In a famous set of early papers that eventually became the definitive early book on high frequency foreign exchange markets [34] show that the impact of high frequency volatility on low frequency volatility is not symmetric.[7] In an interesting paper [78], show that there are patterns in how these volatility asymmetries appear in the data, and many of these are related to specific calendar time periods.

Once one begins talking about variance, then covariances should not be far from the discussion. Modern portfolio theory stresses that the investors should look at covariances of securities in their portfolios as opposed to the variance of any one security. Financial data contains a wealth of information on covariances. I often say that stylized fact number one is that markets often move together, and this might be

[4] See [11] for an early reference to changing volatility patterns. These results created an entire econometric industry of ARCH/GARCH models which are still very important [41] and [13].

[5] Our continuing appreciation of the economic importance of the predictability of variance (and therefore risk) can be seen in modern investment vehicles known as"volatility control" strategies, and recent research such as [83].

[6] It is difficult to tell where this literature starts, but [28] is one of the earliest papers. Similar findings on the extreme persistence in volatility can be found in [4,17,37] and [69].

[7] Much of this work comes from the interdisciplinary work at the foreign exchange research firm, Olsen and Associates, from the 1990's. It was definitely a part of the interdisciplinary Mandelbrot legacy.

for rational or irrational reasons. Econophysics has tackled this problem with several technologies including extensive work on random matrix theory. Another area has been to try break out pockets of correlation, or correlation hierarchies from the data. This has always seemed like a very promising area, and in the current excitement of big data and sophisticated unsupervised learning technology it still seems very appealing.[8] None of these techniques have really caught on in the traditional realm of academic finance. There may be several reasons for this. One is that finance is locked to its own beliefs about the factor structure of asset returns. However, another is that estimating correlations may be a very difficult task statistically, and it is possible that these are changing over time and over time scales.[9]

2.3 Limit order books

In market micro structure and limit-order book modeling the econophysics world has been even more successful. A field with lots of data, and very few established economic theories was ready for more people to enter and plow through its vast amounts quantitative information. Also, distributional analysis, such as power-laws and long memory have proved to be very effective here. Empirical results summarized in [19] show many interesting empirical facts from high frequency trading behavior. Many facts describe the various states of the limit order book which is a mechanism used for trading. Some Mandelbrot like features are repeated as in fat tailed return distributions, and a new related feature was added in the long memory of signed order flow. One example of this would be if you line up buyers as $(+1)$ and sellers as (-1) over time the time series would have very long range autocorrelations as in long memory. A definitive explanation for this important feature is still not agreed on, but it is common in most markets.

Another feature for which the physics community has been actively involved is documenting the shape of the market impact function. This critical property of markets determines how far a certain buying or selling action will move the price as a function of the number of shares traded. If you are a buyer, you push the price up, and a seller you will push it down. The quantity of the price change is the market impact. This is obviously of great interest both to researchers and market participants. It is been determined that this impact function must be close to a square root function of the order size. Testing this result has been one of the more important contributions from the econophysics literature.[10]

This is also an area where there is relatively little economic theory to go by and equilibrium models of the market at the micro second level seem to be pushing some limits to rationality [55]. Here, the physics community has a moved to building highly stylized and very stripped down models for order book dynamics. Models have been built to simulate order book dynamics subject to some reasonable random order flow for some time as in [30]. Several examples of low intelligence order flow from econophysics are [32] and [49]. Similar mechanisms can be seen in [68]. Most of these are developed around either random flows of orders into a market clearing mechanism, or else they use relatively low (zero) intelligence agents as traders.[11] All generate

[8] There are many papers in both these areas, but one of the crucial early papers on hierarchical classification was [81].

[9] See [56] for recent information related to both of these issues, and [65] for information on how correlations patterns change at different time frequencies.

[10] It is interesting that organizing data by universal ratios and scaling invariants has definitely made it into the main stream as in this example [67].

[11] A slightly more intelligent, but not sophisticated, agent is used in the benchmark model of [27].

relatively realistic market dynamics replicating many of the basic stylized facts of high frequency markets. None are completely successful, and there is no accepted model, nor accepted level of sophistication for high frequency traders. However, the low intelligence trader approach has been remarkably useful empirically, and is a very important part of how we think about trading dynamics at high frequencies.

3 Agent-based models and Mandelbrot facts

Agent-based financial models, populated with anywhere from two to thousands of diverse agents and strategies have often taken replication of the facts Mandelbrot discovered in the 1960's as a key goal. This obsession in the agent-based community is matched by almost no theoretical attention in the standard asset pricing world to these same empirical details. The features include the already mentioned fat tailed distributions, and volatility persistence, along with uncorrelated returns which are common in almost all financial series. How well have agent-based models done on these tests? Fortunately, or unfortunately, depending on your perspective, they have been too successful. The world is now loaded with a plethora of agent-based models which replicate the "stylized facts."

The econophysics oriented models fall into two broad categories. First, relatively simple models with stylized structure and relatively few agent types have been built. The emphasis is often on model tractability, and ease of empirical estimation. Some early examples of this are [16,76], and [21]. These are all models which are capable of matching most of the basic features of stock returns, but sometimes require a few minor tweaks to line them up with data.[12] Their dynamics are usually relatively tractable, and various forms of these models can be taken to the data and estimated. For the physics community [16] has probably led to the most activity and extensions as it is built from an Ising model which still feels like a natural structure for information contagion and social interactions [18]. However, this literature still does not seem to empirically favor the Ising model over other approaches. For these simple models, and the basic facts, a definitive empirical test has not been found. We will return to the question of whether these models might be too simple in the next section.

Another interesting early model is the Farmer/Joshi model in [46,47]. This model makes a major move from the other early models in the direction of realistic strategies. The authors are able to still show the impact and dynamics of most of the basic strategies. However, bringing them all together requires computational approaches. The complexity of of this model has made it difficult to estimate relative to some of these other earlier approaches [44].

Another model which overlaps into financial applications is the minority game. The literature on this model is vast. Reference [24] is a good early summary of this large literature. This highly stylized model is an interesting study in game theory, dynamics, and learning, and is a huge and active area for econophysics.[13] From its origin it was suggested to be a good representation for financial market dynamics [23]. This has always seemed like a bit of a stretch. First, it is not a model where price dynamics is part of its core. Prices are usually an add on to the model. Second, it is a model where being a contrarian always wins out. This is not exactly the case when thinking about bubbles and excessive price dynamics in finance. Good strategies in these markets probably follow the herd for a while, but time their departure to avoid the inevitable market crash. At its core, this model is not about this subtle

[12] Examples of these tweaks are in [54] or [50].

[13] The model has roots in the El Farol problem [6] and also shares connections to various earlier models in [85].

aspect of financial market timing. It does have one interesting feature that is a clear contribution of the modeling style in econophysics. Because of its simplicity it is able to enumerate the entire strategy space. For a given history, strategies can be created which cover all possible functions of this history. This is a powerful modeling tool that many more complex models don't have. It is also important to note that this is a model where much of the herd and anti-herd behavior can be viewed as emergent since it does not build in any specific desire to move as a group.[14]

One group of agent-based models that are different from the stylized econophysics models are several early, highly computational, approaches.[15] All these markets push the idea of emergence by leaving the agent strategy set relatively open ended. Agents can learn whatever happens to work, and successful strategies coevolve in a general evolutionary pool of possible forecasts. While intellectually appealing, since little structure is preprogrammed into agents' behaviors, these models can become intractable, and they are also loaded with many parameters. This style of modeling remains distant from the mainstream of econophysics that prefers the much more streamlined models mentioned previously. These models are often capable of replicating the basic empirical features, and in some cases probably do this more effectively than the other simpler models. It remains to be seen whether these models were simply too complex, or maybe slightly ahead of their time.

4 Challenges for agent-based models

Agent-based models are not all that common in the world of economics and finance.[16] If they have such promise, and we are so convinced that heterogeneity matters why aren't these models generally accepted? Some of this has been some general stickiness to behavior in the economics world, but there are definitely other issues that have kept them from broad use.

The first is agent granularity. Most agent-based models depend on some aspect of lumpiness or granularity in the agents. At some point in the system the law of large numbers is breaking down in some subtle way. There is a very good example of this in [39] who take a simple financial model, and then take the limit as the number of agents goes to infinity. The dynamics change dramatically, and most of the interesting features go away.[17] If these models rely on some kind of small numbers issue, which may be important and realistic, we will often need to take the number of agents into account in modeling and estimation. How idiosyncratic noise is added to individual agents is also a problem when they may represent groups rather than true individuals in the model. Also, models where lumpy agents matter may have tricky dynamics which makes analysis either difficult or interesting.

Second, and closely related, is the issue of dynamics. Almost all agent based models generate rich and interesting nonlinear dynamics.[18] These dynamics make these models interesting, and powerful noise generators and amplifiers. Their ability to generate lots of variability in prices from almost no underlying noise is impressive, but it can make estimation difficult. Given that nonlinear dynamical systems can be

14 Another interesting early model which stresses this emergent grouping in a relatively simple way is [89].

15 Some early examples and their respective computational technologies are: classifiers [7], neural networks [10,70], genetic programming [25], and fuzzy logic [88].

16 An interesting exception of an agent-based model in main stream finance is [61].

17 In a slightly more simplified version [3] perform some similar experiments.

18 See [58] for the definitive work on how to formally analyze many of these models which often display low dimensional chaos.

very sensitive to parameter values, they often hand us difficult parameter estimation problems. [19]

Related to estimation is the problem of model comparison. As previously mentioned, when it comes to the Mandelbrot facts, the agent-based world has given us many models which meet the test of feature replication. Therefore, it is probably time to move beyond the basic set of features to others. [74] examines some possibly new facts on the exact dynamics of bubbles [73] looks at a wide array of features at many different time horizons. Finally [20] turns to cross sectional asset pricing relationships. This is a novel approach for the world of agent-based modeling which is often concerned with the price movements of a single asset. I think it is likely that some of these directions in testing will provide tests capable of distinguishing amongst the many models currently in play.

Adaptation and evolutionary dynamics itself can be tricky as well. For finance the agent objectives, or fitness, seem to be well defined in that agents move to better performing strategies. However, what objective is used here? Is this risk adjusted? Should it be? How is this very noisy value estimated from past data? Finance offers some interesting challenges in this area. There actually may be two objectives at work. One coming from direct movements along agents' utility based preferences, and another related to basic wealth growth. The survival of optimal wealth growth portfolios has a long tradition in finance going back to [63], and is directly tied to some basic ideas from information theory. Financial markets that allow for wealth adjustments along with some form of response to recent performance implicitly are operating with two evolutionary objectives [71]. This combination of dynamics is still not a well understood area.[20] Even just one fitness measure is tricky since the market demands will often depend critically on its numerical values. This means the magnitude of a relative comparison across two strategies can be very important since this will feed into a gradient flow of agents between the strategies. As researchers we may be very sure of making relative comparisons between these strategies, but the quantitative comparisons necessary to determine agent dynamics is probably something we are much less sure about.

Many agent based models stress their adaptive dynamics. Sometimes this comes from evolutionary models, and sometimes models for social interactions. Whichever mechanism is used, little modeling effort is made at the level of the agent forecasting or behavioral rules. This is even true in places where there is data available to calibrate these rules. Researchers work on models with rich and interesting dynamics across agents who may simply be optimists or pessimists, buyers or sellers, trend followers, or fundamentalists. These very simple agent behaviors are a weak link in this modeling chain. There is more information out there, and we can do better. Some of this data is not pretty, but it can be brought in to build a better agents. A recent example of this is [1] where the authors use investor forecasting data to calibrate learning memory lengths for a set of trading agents. References [64] and [51] are other examples of agent calibration. Another method is used in [72] which replaces what are often ad hoc rules for fundamental traders with strategies based on common price/dividend ratio regressions. Beyond information from the field, there is also a vast amount of information from controlled laboratory experiments that can be used as well for calibrating agents. Reference [57] is a good survey of this. Also, see [59] for some very interesting connections between experimental results and computational simulations.

[19] See [14] for an example where the difficulty of estimating some critical model parameters is very clear. Another example of model estimation is [2]. An early simple experiment on the difficulties of estimating a chaotic model using traditional methods is [53].

[20] See [5] and [15] for some early directions. Also, a very complete survey of this area is in [43].

Getting more detailed agent calibration information is an obvious forward step for agent-based analysis.

5 Sumary and future

As I've mentioned, replicating Mandelbrot's puzzles is something agent-based models have been very successful at. Unfortunately, this is not going to gain them acceptance into the more traditional economics modeling toolkit. What will eventually be needed is some model that dramatically changes how we think about a problem, and hopefully delivers some new and previously not considered policy advice. For the most part this has not happened yet.

Sometimes agent-based models may make policy recommendations which are too far reaching and radical to be believed. This might be the case in [36] who hinted at some of the dangers of changing the way prices were quoted on NASDAQ. The model is relatively complex, and the change from eighths to decimal trading seemed to be destined to happen regardless of the authors' simulated markets showing the possible appearance of new parasitic strategies under the different institutional arrangement. Another example coming from politics is in [66]. They use an agent-based model along with some intuition coming from simulated annealing to suggest certain optimal voting systems for local politics. I don't think much of this has been implemented in real election systems. In both cases we may have met the intellectual bar of an interesting and important model, but the world was not ready for its message.

Agent-based models are often asked to do what we ask from traditional models. This usually is in some form of out of sample prediction, or prediction subject to some policy intervention or treatment. They may not be the best models for doing this. It is possible that these models will never be any better (or might be worse) than more traditional models at predicting financial crises. Given that they need to calibrate agents to data pulled from the rare periods when markets are under stress, and also that they involve interesting, but challenging, nonlinear dynamics, predicting extreme events may not be something they can do well. This should not be held out as a test for their usefulness.[21] In their defense, agent-based models may be perfect counters to the Lucas Critique [75] in which empirically derived policy recommendations can go terribly wrong as agent behavior changes under a new regime. Well designed agent-based models would learn around the new policies predicting the new outcome. This is interesting, but is admittedly a stretch given our current technologies.

More important results will come from challenging more traditional approaches. A notion of breaking standard models might be a more fruitful approach. An interesting engineering example of this is the Millennium Bridge [12]. Building bridges involves taking well tested models and taking them "out of sample." The laws of physics and mechanics hold well as civil engineers have used these successfully for thousands of years. The Millennium Bridge was built based on our usual bridge making technologies, but it failed to take into account the endogenous dynamics caused by pedestrians walking across the swaying bridge. An agent-based model might have captured this phenomenon and offered some useful design interventions. Another example is the concept of "slower is faster" [52] in which putting barriers into a flow pattern can have the impact of speeding up pass through. It would be interesting to find big, nonintuitive features such as these for financial markets.

Econophysics has followed in the legacy left by Benoit Mandelbrot. Often defiantly operating outside the bounds of economics and finance, it has made contributions

[21] Though I remain cautious about building early warning indicators, some recent research suggests this may be feasible. See [38] for measures of liquidity near a crash, and [87] and [9], for predictions using banking network information.

that have added to our knowledge about financial market operations. On the empirical side, by concentrating more on big stylized facts, it has revealed many patterns in financial data. It is important to single out market microstructure, and in particular limit order books, as a major success for the field. As for agent-based modeling, econophysics has probably fared no better, or worse, than the fringe of economics involved in this modeling. It has purposefully concentrated on highly tractable, very stripped down models of behavior. My fear is that they may be too stripped down for what they are trying to do. As our data and computing resources continue to expand, I remain optimistic about this research area operating between the physical and social sciences.

References

1. K. Adam, A. Marcet, J. Monetary Econ. (2016) forthcoming.
2. S. Alfarano, T. Lux, F. Wagner, Comput. Econ. **26**, 19 (2005)
3. V. Alfi, M. Cristelli, L. Pietronero, A. Zaccaria, Eur. Phys. J. B **67**, 399 (2009)
4. T.G. Andersen, T. Bollerslev, F.X. Diebold, P. Labys, Econometrica **96**, 579 (2003)
5. M. Anufriev, P. Dindo, J. Econ. Behavior and Organization **73**, 327 (2010)
6. W.B. Arthur, Am. Econ. Rev. **84**, 406 (1994)
7. W.B. Arthur, J. Holland, B. LeBaron, R. Palmer, P. Tayler, Asset Pricing Under Endogenous Expectations in an Artificial Stock Market, in W.B. Arthur, S. Durlauf, D. Lane, editors, *The Economy as an Evolving Complex System II* (Addison-Wesley, Reading, MA, 1997), pp. 15–44
8. O.E. Barndorff-Nielsen, N. Shephard, J. Fin. Econometrics **2**, 1 (2004)
9. S. Battiston, J.D. Farmer, A. Flache, D. Garlaschelli, A.G. Haldane, H. Heesterbeek, C. Hommes, C. Jaeger, R. May, M. Scheffer, Science **351**, 818 (2016)
10. A. Beltratti, S. Margarita, P. Terna, *Neural Networks for Economic and Financial Modeling* (International Thomson Computer Press, London, UK, 1996)
11. F. Black, Studies of Stock Price Volatility Changes, in *Proceedings of the American Statistical Association, Business and Economics Statistics Section, 1976*, pp. 177–181
12. A. Blekherman, Int. J. Bridge Eng. **3**, 1 (2015)
13. T. Bollerslev, Rev. Econ. Stat. **69**, 542 (1987)
14. H.P. Boswijk, C.H. Hommes, S. Manzan, J. Econ. Dyn. Control **31**, 1938 (2007)
15. G. Bottazzi, P. Dindo, J. Econ. Dyn. Control **48**, 121 (2014)
16. J.P. Bouchaud, R. Cont, Macroecon. Dyn. **4**, 170 (2000)
17. J.P. Bouchaud, I. Giardina, M. Mezard, Quant. Fin. **1**, 212 (2001)
18. J.-P. Bouchaud, J. Stat. Phys. **151**, 567 (2013)
19. J.-P. Bouchaud, J.D. Farmer, F. Lillo, How Markets Slowly Digest Changes in Supply and demand, in *Handbook of Financial Markets: Dynamics and Evolution*, (North-Holland, 2009), pp. 57–160
20. J.-P. Bouchaud, P. Krueger, A. Landler, D. Thesmar, Stick Expectations and Stock Market Anomalies, Technical report, HEC, Paris, 2016
21. W.A. Brock, C.H. Hommes, J. Econ. Dyn. Control **22**, 1235 (1998)
22. A. Chakraborti, I.M. Toke, M. Patriarca, F. Abergel, Quant. Fin. **11**, 991 (2011)
23. D. Challet, A. Chessa, M. Marsili, Y.-C. Zhang, Quant. Fin. **1**, 168 (2001)
24. D. Challet, M. Marsili, Y.C. Zhang, *Minority Games* (Oxford University Press, 2004)
25. S.-H. Chen, C.-H. Yeh, J. Econ. Dyn. Control **25**, 363 (2001)
26. C. Chiarella, R. Dieci, X.-Z. He, Heterogeneity, Market Mechanisms, and Asset Price Dynamics, in T. Hens and K. R. Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution* (Elsevier, USA, 2009), pp. 277–344
27. C. Chiarella, G. Iori, Quant. Fin. **2**, 346 (2002)
28. P. Cizeau, Y. Liu, M. Meyer, C.K. Peng, H.E. Stanley, Physica A **245**, 441 (1997)
29. A. Clauset, C.R. Shalizi, M.E.J. Newman, SIAM Rev. **51**, 661 (2009)
30. K.J. Cohen, S.F. Maier, R.A. Schwartz, D.K. Whitcomb, Simulation **41**, 181 (1983)

31. R. Cont, IEEE Signal Processing Magazine **28**, 16 (2011)
32. R. Cont, S. Stoikov, R. Talreja, Operations Res. **58**, 549 (2010)
33. P. Cootner, editor, *The Random Character of Stock Market Prices* (MIT Press, Cambridge, 1964)
34. M.M. Dacorogna, R. Gencay, U.A. Muller, R.B. Olsen, O.V. Pictet, *An Introduction to High-Frequency Finance* (Academic Press, San Diego, CA, 2001)
35. J. Danielsson, K. James, M. Valenzuela, I. Zer, J. Fin. Stability (2016) forthcoming
36. V. Darley, A.V. Outkin, *A NASDAQ Market Simulation*, Complex Systems and Interdisciplinary Science (World Scientific, 2007)
37. Z. Ding, C.W.J. Granger, R.F. Engle, J. Empir. Fin. **1**, 83 (1993)
38. J. Donier, J.-P. Bouchaud, PLOS One **10**, e0139356 (2015)
39. E. Egenter, T. Lux, D. Stauffer, Physica A **268**, 250 (1999)
40. P. Embrechts, C. Kluppelberg, T. Mikosch, *Modeling Extremal Events for Insurance and Finance* (Springer-Verlag, 1997)
41. R.F. Engle, Econometrica **50**, 987 (1982)
42. J. Estrada, J. Investing **17**, 20 (2008)
43. I.V. Evstigneev, T. Hens, K.R.S.-Hoppe, Evolutionary finance, in Thorsten Hens and Klaus Reiner Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution*, Handbooks in Finance (North-Holland, Amsterdam, the Netherlands, 2009), pp. 509–564
44. A. Fabretti, J. Econ. Interactions and Coordination **8**, 277 (2013)
45. E.F. Fama, J. Business **36**, 420 (1963)
46. J.D. Farmer, Industrial and Corporate Change **11**, 895 (2002)
47. J.D. Farmer, S. Joshi, J. Econ. Behavior and Organization **49**, 149 (2002)
48. J.D. Farmer, L. Gillemot, F. Lillo, S. Mike, A. Sen, Quant. Fin. **4**, 383 (2004)
49. J.D. Farmer, P. Patelli, I. Zovko, Proc. Natl. Acad. Sci. U.S.A. **102**, 2254 (2005)
50. A. Gaunersdorfer, C. Hommes, A Nonlinear Structural Model for Volatility Clustering, in A. Kirman and G. Teyssiere, editors, *Micro Economic Models for Long Memory in Economics* (Springer-Verlag, 2007), pp. 265–288
51. J. Geanakoplos, R. Axtell, D. Farmer, P. Howwitt, B. Conlee, J. Goldstein, M. Hendrey, N. Palmer, C.-Yi Yang, Am. Econ. Rev. **102**, 1 (2012)
52. C. Gershenson, D. Helbing, Compexity **21**, 9 (2015)
53. J. Geweke, Inferences and Forecasting for Deterministic Non-linear Time Series Observed With Measurement Error, in R.H. Day and P. Chen, editors, *Nonlinear Dynamics and Evolutionary Dynamics* (Oxford University Press, 1993)
54. F. Ghoulmie, R. Cont, J.-P. Nadal, J. Phys.: Condens. Matter **17**, 1259 (2005)
55. R. Goettler, C.A. Parlour, U. Rajan. J. Fin. **60**, 2149 (2005)
56. C.R. Harvey, Y. Liu, H. Zhu, Rev. Fin. Studies **29**, 5 (2016)
57. C.H. Hommes, J. Econ. Dyn. Control **35**, 1 (2011)
58. C.H. Hommes, *Behavioural Rationality and Heterogeneous Expectations in Complex Economic Systems* (Cambridge University Press, Cambridge, UK, 2013)
59. C.H. Hommes, T. Lux, Macroeconomic Dynamics **18**, 373 (2013)
60. C.H. Hommes, F. Wagener, Complex Evolutionary Systems in Behavioral Finance, in Thorsten Hens and Klaus Reiner Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution* (North-Holland, 2009), pp. 217–276
61. B.I. Jacobs, K.N. Levy, H.M. Markowitz, Fin. Analyst. J. **66**, 42 (2010)
62. B. Kelly, H. Jiang, Rev. Fin. Studies **27**, 2841 (2014)
63. J.L. Kelly, Bell Syst. Tech. J. **35**, 917 (1956)
64. A.E. Khandani, A.W. Lo, R.C. Merton, J. Fin. Econ. **108**, 29 (2013)
65. W.B. Kinlaw, M. Kritzman, D. Turkington, The Divergence of High and Low-frequency Estimation: Implications for Performance Measurement, Technical report, State Street Global Exchange, 2014
66. K. Kollman, J. Miller, S.E. Page, Am. Econ. Rev. **87**, 977 (1997)
67. A.S. Kyle, A.A. Obizhaeva, Market Microstructure Invariants: Empirical Hypothesis, Technical Report, University of Maryland, 2016
68. D. Ladley, KR. Schenk-Hoppe, J. Econ. Dyn. Control **33**, 817 (2009)

69. B. LeBaron, Quant. Finance **1**, 621 (2001)
70. B. LeBaron, Agent-based Financial Markets: Matching Stylized Facts With Style, in D. Colander, editor, *Post Walrasian Macroeconomics* (Cambridge University Press, 2006), pp. 221–235
71. B. LeBaron, East. Econ. J. **37**, 35 (2011)
72. B. LeBaron, J. Econ. Behav. Organ. **83**, 424 (2012)
73. B. LeBaron, Heterogeneous Agents and Long Horizon Features of Asset Prices, Technical report, Brandeis International Business School, 2013
74. M. Leiss, H.H. Nax, D. Sornette, J. Econ. Dyn. Control **55**, 1 (2015)
75. R.E. Jr. Lucas, Carnegie-Rochester Conference Series on Public Policy **1**, 19 (1976)
76. T. Lux, M. Marchesi, Nature **397**, 493 (1999)
77. T. Lux, Stochastic Behavioral Asset Pricing Models and the Stylized Facts, in Thorsten Hens and Klaus Reiner Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution* (North-Holland, 2009), pp. 161–215
78. P. Lynch, G. Zumbach, Quant. Fin. **3**, 320 (2003)
79. B.B. Mandelbrot, J. Business **36**, 394 (1963)
80. R.N. Mantegna, H.E. Stanley, Nature **376**, 46 (1996)
81. R.N. Mantegna, Eur. Phys. J. B **11**, 193 (1999)
82. R.C. Merton, J. Fin. Econ. **3**, 125 (1976)
83. A. Moreira, T. Muir, Volatility Managed Portfolios. Technical Report, Yale School of Management, 2015
84. P. Savor, M. Wilson, J. Fin. Econ. **113**, 171 (2014)
85. T. Schelling, *Micromotives and Macrobehavior* (Norton, New York, NY, 1978)
86. D. Sornette, Rep. Prog. Phys. **77**, 062001 (2014)
87. T. Squartini, I. van Lelyveld, D. Garlaschelli, Sci. Rep. **3**, 3357 (2013)
88. N.S.P. Tay, S.C. Linn, J. Econ. Dyn. Control **25**, 321 (2001)
89. M. Youssefmir, B.A. Huberman, J. Econ. Behav. Organ. **32**, 101 (1997)