Defining efficiency in heterogeneous markets

Michel Dacorogna, Ulrich Müller, Richard Olsen and *Olivier Pictet* offer a new definition of efficient markets that takes into account heterogeneity and varying time scales.

In the last few years, research in finance has increasingly focused on the behaviour of high-frequency data. Some results of this research provide clear evidence that movements in foreign exchange (FX) rates and in the prices of other financial assets, for short to medium-term horizons, are, to some extent, predictable. This is substantiated by a positive forecast quality and high real-time trading model returns. More generally, financial returns on assets are seen to depart substantially from the random walk model—and market participants are predicting them with some success.

Where does this sustained predictability originate from? Are the real-time trading models successful in capturing the inefficiencies of the FX market, for instance? Since this market is widely held to be the most efficient of the financial markets, does such success conflict with the theory of efficient markets (which precludes the ability to forecast and denies the existence of profitable trading models)? Or should we conclude that markets are inefficient?

We prefer to adapt our theory of financial markets to the reality, in which markets are very efficient but in a newly defined way. This essay aims to explain why and how markets can be highly efficient and, at the same time, to some extent predictable.

There are several reasons for this, all associated with market dynamics. We want to put in perspective the current theory of efficiency and suggest moving beyond it. This is a major challenge faced by those working in the theory of finance, which include the 'behavioural finance' movement around Robert Shiller, parts of the econophysics community and many others, who see the need to find ways of moving from a rather static definition to a more dynamic one (exactly the aims of the workshop 'Beyond Equilibrium and Market Efficiency', where this paper was first presented).

Definition of efficient markets

In conventional economics, markets are assumed to be efficient if all available information is reflected in current market prices. Economists have developed three types of test for efficiency. The weak-form tests investigate whether market prices actually reflect all available information. The semi-strong tests are based on so-called event studies, where the degree of market reaction to news announcements is analysed. Finally, the strong-form tests analyse whether specific investors or groups have private information from which to take advantage. Most studies conclude that the major financial markets are efficient and that all information is reflected in current prices. However, their conclusions have been bogged down by methodological questions, in particular relating to whether observed departures from market efficiency are due to genuine market inefficiency or whether there is a deficiency in the market pricing model used to compare actual with theoretical prices.

The inference that trading models cannot generate excess return in an efficient market is based on the assumption that all investors act according to the rational expectation model. If this assumption is wrong, we must also question the conclusion (that forecasting is impossible). The assumption of rational expectations has been called into question on various platforms and the idea of heterogeneous expectations has become of increasing interest.

Shiller, for example, argues that most participants in the stock market are not 'smart investors' (following the rational expectation model) but rather follow trends and fashions. The modelling of 'noise trader' has become a central subject of research in market microstructure models. On the FX market, there is much investigation of 'speculative bubbles' and the influence of technical analysis on the dealer's strategy.

Attention has also been caught by the possibility of time-varying expectations, which is closer to our view of the market. Variation in expected returns over time poses a challenge for asset-pricing theory because it requires an explicit dynamic theory in contrast to the traditional static capital asset pricing model.

In summary, the conclusion that financial asset prices are not predictable is based on three assumptions: that market prices reflect all the information available, that

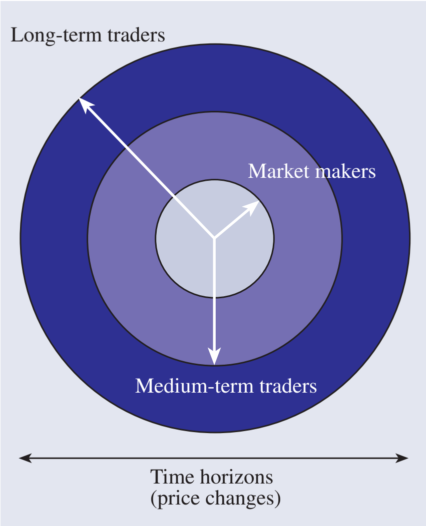

Figure 1. Different types of trader operate over different time scales.

news and events hitting the market are normally distributed and that the market is composed of homogeneous agents. The two first assumptions are reasonable starting points for the definition. The third assumption poses a real problem. It is clear that all market agents in fact have bounded rationality. They cannot be omniscient and do not all enjoy the same freedom of action and access to the markets. Recent works by Kurz and by Gouree and Hommes present new theoretical models to tackle this problem. Introducing the heterogeneity of agents can give rise to very interesting nonlinear effects in the models. They show that many of the price fluctuations can be explained by endogenous effects. Similar conclusions are reached by Farmer and Lo in their discussion of market efficiency. They base their analysis on a comparison with the evolution of ecological systems.

Dynamic markets and relativistic effects

For a model based on efficient markets in which economic agents are entities acting according to the rational expectation strategy, any differences in planning horizons, frequency of trading or institutional constraints are neglected.

The assumption of rational expectations has been called into question on various platforms and the idea of heterogeneous expectations has become of increasing interest.

However, there is substantial empirical evidence that investors have heterogeneous expectations. Surveys on the forecasts of participants in the FX market reflect the wide dispersion of expectation at any point in time. The huge volume of FX trading (over $1500 bn a day across the main financial centres) is another indication that reinforces this idea: it takes differences in expectation among market participants to explain why they trade.

The concept of heterogeneous markets, supported by empirical facts, is the most elegant way to reconcile market efficiency with the stylized facts. Recently, Lux and Marchesi developed simulation models of financial markets that include agents with different strategies (fundamentalists and chartists). They were able to show that this model can reproduce most of the empirical regularities (fat tails, long memory and scaling law) even though they use normally distributed news in their simulations.

There are many ways to describe these heterogeneous expectations. We believe that the most promising approach is to differentiate the expectations according to their time dimension because we consider the *different time scales* of the market participants the key characteristic of the market. Some trade short term, others have long-term horizons. Market makers are at the short-term end of the scale and central banks at the long-term end. Contrary to the usual assumption, there is no privileged time scale in the market. The interaction of components with different time scales gives rise to characteristically relativistic effects (N.B. we use the term 'relativistic' to express the dynamic interaction between different market components relative to each other—rather than relative to the news that has impacted the market). Examples of such effects include certain properties of volatility clusters, trend persistence, and the lag between interest rate adjustment and FX rate adjustment.

The latter is a good example of what conventional theory considers an inefficiency while we see it as an effect arising from the different time scales involved in the market. To take advantage of the lag in adjustment between interest rate and exchange rate moves, an investor needs to tie up his money for months or even years. This is a very long time for an FX trader. Some investors will thus tend to ignore these profit opportunities while others invest in them. The combination of all of these effects ultimately enables the construction of successful forecasting and trading models.

In long time intervals, market price changes are 'flatter' and have fewer relevant movements (trend changes) than in short-term intervals. The higher the resolution and the smaller the intervals, the larger the number of relevant price movements. The long and the short-term traders thus have different trading opportunities: the shorter the trading horizon, the greater the opportunity set. A market participant's response to outside events should always be viewed as relative to their intrinsic opportunity set. Economic decision makers, such as traders, treasurers and central bankers, interpret the same information differently.

This variation in perspective has the effect that specific price movements cannot lead to a uniform reaction; rather, they result in individual reactions of different components. In turn, these reactions give rise to secondary reactions, with the different components reacting to their respective initial response. Watching the intraday price movements, one clearly sees the sequences of secondary reactions triggered by the initial events. The existence of different trading strategies in the market was also put forward by Müller *et al* to explain asymmetry in the information flow at different frequencies. There is no competing explanation of this significant effect which states that coarse volatility (where returns are measured over large time intervals) predicts fine volatility (where returns are measured over small intervals) better than the other way around. LeBaron shows that introducing agents with different time horizons in his market model gives rise to long-term clustering effects in the resulting price volatility (or autoregressive heteroskedasticity).

The delay with which the secondary reactions unfold is called the *relaxation time*. If diverse components with different time scales interact in the market, there is typically a mixture of long and short relaxation times following the impact of outside events. If different relaxation times are combined, the resulting autocorrelation decays hyperbolically or almost hyperbolically. This is a natural explanation of the long memory effects detected in financial markets and has been studied already by Dacorogna and others.

There is yet another phenomenon to consider, which originates from the fact that financial markets are spread worldwide. Economic and political news and trading activity are not stationary. They move around the world in a 24-hour cycle. The price data of foreign exchange rates reflects this, in terms of a 24-hour seasonality in market volatility. This cycle implies that market reactions to an event cannot be simultaneous and that there are distinct relaxation times following an event. Geographical components related to the business hours of the different trading centres must be added to the time components and the interaction of these geographical components leads to behaviours such as the 'heat wave' effect proposed by Engle *et al.*

Impact of the new technology

The realization that there is value in the data for defining investment strategies has brought to life many new firms specializing in modelling financial markets and in providing trading advice on the basis of technical models. The question is, of course, will the impact of the new technology be a passing phenomenon or will it have a long-term effect? As the relativistic phenomenon arises from the interaction of components with different time scales, it will remain appropriate as long as heterogeneous expectations continue to exist in the market. The interaction process may become more complex, but it cannot disappear.

New technologies enable users to identify additional trading opportunities to increase their profits. This quickens their pace of trading and contributes to higher market volume and liquidity. The improved liquidity lowers the spreads between bid and ask prices. Lower spreads imply lower transaction costs, which in turn increase the opportunity horizon for profitable trading. The new technology introduces a shift in perspective, with components starting to focus on more numerous short-term time intervals.

As components become increasingly short term in their focus, the spectrum of short-term components increases. This has the effect that relative differences among components become more significant and the relativistic effects more pronounced. Contrary to accepted notions which assume that sufficient buying power can 'trade away' any phenomenon, the increased buying power will have the overall effect of enhancing the relativistic effects. Thus the very basis of our ability to forecast and build profitable trading models will be enhanced. It should be noted, however, that the reaction patterns will become increasingly diversified and therefore more complex and so the speed of adjustment will increase, requiring ever more sophisticated models.

Zero-sum game or *perpetuum mobile*?

Conventional thought has it that financial markets must be a zero-sum game. This is true if we take a static view. In reality, financial markets are dynamic and highly complex. Markets are a platform for components to exploit the diversity of interests. They can match their opposing objectives when one component buys and another component sells. The lower the friction, the easier a counterpart for a particular transaction is found and the larger, therefore, is the particular component's opportunity set. The ability to proceed with a particular transaction means that the flexibility of the respective components is increased and their profit potential improved.

The new technology fosters the ability of the market to provide an environment for the generation of wealth. Interaction within the market gives rise to relativistic effects and relaxation times. To the extent that these relativistic effects are understood and incorporated into forecasting and trading model technology, market participants have the opportunity to generate additional profit or limit their losses. In our terminology, the profit generated is energy extracted from the market. Improved efficacy of component interaction generates additional energy and reduces the friction associated with buying and selling within the market. The process may be compared to the search for more efficient engines in the automobile industry where everybody gains from it in the long term.

Have we achieved a *perpetuum mobile*? The answer is clearly no. Like any other technological innovation, the new technology does not generate energy from nothing, but it does take advantage of the energy potential existing in the financial markets. By offering a service to the economic agents, financial markets are not closed systems but have a permanent input of money. This makes them highly open systems in terms of energy. Besides, a lot of resources have been put into the new technology in the form of extensive research, development work and hardware to treat the information. Numerous studies have shown that simple trading rules do not work in efficient markets. Only elaborate treatment of the data allows the identification of profitable trading rules. This treatment is not free: it has a price. Moreover, as the relativistic reaction patterns become increasingly diversified, research and development efforts will have to increase to keep up with ever-changing nonlinear patterns.

The conventional definition

As the markets comprise a diversity of components, there are different relaxation times, arising from the underlying relativistic effects between different components. It follows that the weak form of efficiency coupled with the rational expectation model cannot be attained. Because of the presence of different time components with heterogeneous expectations, current market prices cannot reflect all available information. Instead the price discovery mechanism follows a dynamic 'error correction model' where the successive reactions to an event unfold in the price. Why, then, did this not show up more clearly in previous scientific investigations? Reasons include the following:

- high-frequency data is a prerequisite for the empirical investigation of relativistic phenomena;
- extensive computing power is needed to show the predictability in financial markets. Access to reasonably priced computing power has become available only recently;

The new technology introduces a shift in perspective, with components starting to focus on more numerous short-term time intervals.

- only in the last few decades has an increasing awareness for dynamic and nonlinear processes been gained. Such an awareness is crucial for the study of relativistic effects.

The presumption of conventional economics that forecasting is impossible by definition has had a powerful impact on research into market efficiency. Economists focused on structural studies which were hamstrung by a lack of high-frequency data and theoretical shortcomings. Little academic research has been invested in trying to predict shorter-term price movements and build successful trading models.

An improved definition of efficient markets

Although the current definition of efficient markets has shortcomings, we do not think that the concept should be abandoned; rather, it should be adapted to the new findings. It is important to find a good measure of how well a market operates.

From a dynamic perspective, the notion of reduced friction should be central to the notion of efficiency. We consider an efficient market to be a market where all market information must be available to the decision makers and where there must be participants with different time scales and heterogeneous expectations trading with each other to ensure a minimum of friction in the transaction costs.

A quantitative measure of efficiency might be derived from the bid–ask spreads (the spreads between real bid and ask prices being more appropriate than the nominal spreads quoted in information systems). Spreads are not only a measure of 'friction', they also contain a risk component. The volatility or, more precisely, the probability of extreme price changes within short time intervals, should be considered together with the spread, in the quantitative measure of market efficiency to be proposed. We are confident that in the years to come this definition will prevail and we will find precise measures of efficiency as is the case in thermodynamics and engineering.

Further reading

Almeida A *et al* 1998 The effects of macroeconomic news on high frequency exchange rate behavior *J. Financial Quantitative Anal.* **33** 383–408

Andersen T G and Bollerslev T 1997 Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns *J. Finance* **52** 975–1005

Baillie R T 1996 Long memory processes and fractional integration in econometrics *J. Econometrics* **73** 5–59

Bank for International Settlements 1999 Central bank survey of foreign exchange and derivatives market market activity 1998 *Monetary and Economic Department, Basle, Document*

Bekaert G and Hodrick R J 1992 Characterizing predictable components in excess returns on equity and foreign exchange markets *J. Finance* **47** 467–509

Dacorogna M M *et al* 1992 Short-term forecasting models of foreign exchange rates *Presentation at the IBM Summer Research Institute in Oberlech Austria on Advanced Applications in Finance, Investment and Banking, 27–31 July 1992* (Zürich: Olsen & Associates)

Dacorogna M M *et al* 1993 A geographical model for the daily and weekly seasonal volatility in the FX market *J. Int. Money Finance* **12** 413–38

Ding Z *et al* 1993 A long memory property of stock market returns and a new model *J. Empirical Finance* **1** 83–106

Engle R F *et al* 1990 Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market *Econometrica* **58** 525–42

Fama E F 1970 Efficient capital markets: a review of theory and empirical work *J. Finance* **25** 383–417

Fama E F 1991 Efficient capital markets: II *J. Finance* **46** 1575–617

Farmer J D 1998 Market force, ecology, and evolution *Santa Fe Institute preprint* 1–66

Farmer J D and Lo A W 1999 Frontiers of finance: evolution and efficient markets *Proc. Natl Acad. Sci. USA* **96** 9991–92

Franke G and Hess D 1997 Information diffusion in electronic and floor trading *University of Konstanz Discussion Paper*

Frankel J A and Froot K A 1990 Chartists, fundamentalists, and trading in the foreign exchange market *Am. Econ. Rev.* **80** 181–85

Gencay *et al* 1998 Real-time trading models and the statistical properties of foreign exchange rates *RAG Internal Document* (Zürich: Olsen & Associates)

Goodhart C A E and Curcio R 1991 The clustering of bid/ask prices and the spread in the foreign exchange market *LSE Financial Market Group Discussion Paper Series* **110** 1–15

Gouree J K and Hommes C H 2000 Heterogeneous beliefs and the nonlinear cobweb model *J. Econ. Dynamics Control* **24** 761–98

Kurz M 1994 Endogenous economic fluctuations, studies in the theory of rational beliefs *Studies in Economic Theory* vol 6 (Berlin: Springer)

LeBaron B 2000 Evolution and time horizons in an agent based stock market *Working paper* available at www.unet.brandeis.edu/~blebaron

Lux T and Marchesi M 1999 Scaling and criticality in a stochastic multi-agent model of a financial market? *Nature* **397** 498–500

Lux T and Marchesi M 2000 Volatility clustering in financial markets: a microsimulation of interacting agents? *Int. J. Theor. Appl. Finance* **4** 1–20

Müller U *et al* 1990 Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis *J. Banking Finance* **14** 1189–208

Müller U *et al* 1993 Fractals and intrinsic time—a challenge to econometricians *Invited presentation at the XXXIXth International AEA Conference on Real Time Econometrics, 14–15 October 1993 (Luxembourg)* and the *4th International PASE Workshop, 22–26 November 1993 (Ascona, Switzerland)*; also in 1994 *Erfolgreiche Zinsprognose* ed B Lüthje (Zürich: Olsen & Associates)

Müller U *et al* 1997 Volatilities of different time resolutions—analyzing the dynamics of market components *J. Empirical Finance* **4** 213–39

Pictet O V *et al* 1992 Real-time trading models for foreign exchange rates *Neural Network World* **2** 713–44

Shiller R J 1989 *Market Volatility* (Cambridge, MA: MIT Press)

Shiller R J 2000 *Irrational Exuberance* (Princeton, NJ: Princeton University Press)

Michel Dacorogna is based at Zurich Re, Zürich; Ulrich Müller and Richard Olsen are at Olsen & Associates, Zürich; Olivier Pictet is at Dynamic Asset Management, Geneva.