Time Scales, Agents, and Empirical Finance

Blake LeBaron *

Brandeis University

May 2006

Abstract

This article looks at empirical features and how they span multiple time scales in financial time series. I claim that these connections across different frequencies are important in the world of finance, and are related to how trading agents interact with an environment that is, to some extent, their own creation. Lack of a clear time scale allows many heterogeneous beliefs to flourish, and it is the coexistence of these beliefs which helps to generate features across many times scales in a kind of symbiotic relationship.

*Abram L. and Thelma Sachar Chair of International Economics, International Business School, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a faculty research fellow at the National Bureau of Economic Research.

Macroeconomics and finance are often treated separately in economic research. Tinbergen was one of the early pioneers for suggesting that financial markets need to be integrated with analysis of the overall macro economy. He also has several interesting papers which are connected to some of what I will be discussing here. First, as someone interested in dynamics he was well aware of the importance of time scales, or planning horizons, and how this affects economic agents, (Tinbergen 1933). The importance of different horizons will be a common theme in this article. It is also interesting that his models of the stock market considered agents looking at aspects of recent past performance as a possible tool for current price determination, (Tinbergen 1939). This is commonly referred to as a "momentum" effect in finance, and appears in my work as trend following technical trader behavior.

This short article will quickly summarize some of the more unusual empirical features in finance. A common aspect of these is that they still remain difficult to explain in any theoretical model. I will summarize some recent work which tries to build models which replicate these features. A common argument through all of this is that how one considers time is critical, and up to now we have generally taken too simple an approach, by concentrating models on specific frequencies in the data. I will argue that a more unified approach, spanning the highest to the lowest frequencies may be necessary to completely understand the dynamics of financial markets.

1 Multiple time scales, stationarity, and long memory

Throughout this article there is a common theme that these three issues are an important, underexplored area of financial time series. They are critical both to understanding many empirical features as well as constructing models to replicate them. Most empirical and theoretical work in economics makes assumptions which mitigate any problems coming from all of these. In empirical modeling we often concentrate on a single time frame (daily, monthly, annual).[1] It is also often useful to both assume our time series are stationary, and that the economic agents who populate our models share this opinion. Problems of time scales and stationarity get mixed together when one considers long memory processes. Long memory, or fractionally integrated, processes have shocks that die off at a hyperbolic rather than an exponential rate. We often assume this away by staying with the standard toolbox of stochastic processes.

There is a large overlap both empirically, and theoretically between all three of these issues. Time series that contain nonstationary changes in policy or regime can appear as long memory to certain tests. The

[1] Early examples that try to incorporate multiple scales are described in (Dacorogna, Gencay, Muller, Olsen & Pictet 2001). Another example is the MIDAS model used in (Ghysels, Santa-Clara & Valkanov forthcoming 2006).

addition of time series with dynamics operating on different time scales can also appear as long memory.[2] It is also possible that periodic policy changes may confuse economic decision makers as to the appropriate time scales to estimate econometric models. For example, forecasting the long run returns to stocks and bonds would be a lot easier if we were completely convinced that the last 100 years was a representative draw from the true stationary distribution.

This triple set of problems pervades all of macroeconomics, and not just finance. However, with its wealth of data ranging in frequencies from seconds to years, finance is probably one of the best places to attack these issues. I will argue that this is a direction the field is moving in, and is a trajectory which may remain important for many years into the future.

2 Empirical features of financial time series

The set of features described here is pretty general, and all have been seen in many different asset returns. This would include equity markets, foreign exchange markets, money markets, and commodity markets. It is important that these general stylized facts are not specific to any country, time period, or asset class.

I will quickly mention some of the most well known features since they are described in many places, and only indirectly related to the time scale themes of this article. First, financial markets are difficult to forecast. Prices follow something close to a martingale in log first differences,

$$\log(p_{t+1}) = \log(p_t) + r_{t+1}, \tag{1}$$

where r_{t+1} refers to the return from t to $t+1$. We know that for all asset markets the martingale nature appears in that r_{t+1} is relatively difficult to forecast using time t information. In the remainder of this article, I will follow usual practice and discuss the dynamics of r_t, and not p_t.

A second empirical issue coming from traditional financial models is that we have a difficult time reasonably connecting risk and return, both in the cross section, and time series dimension. The Capital Asset Pricing Model (CAPM) was a simple, elegant approach to risk and return, but the most basic form of the CAPM is not successful in explaining cross sectional stock returns.[3]

Both of these are standard results, and well known to almost anyone in economics and finance. I will

[2] See (Granger 1980) or (LeBaron 2001) for examples.

[3] It is also connected to the consumption CAPM, the APT, and several other models. Variants of these continue to be "tweaked" with more or less success, but no model has stood out as a clear winner. The basic message of (Roll 1987) still holds today that we really are not very good at explaining much in financial data.

now move to some of the lesser known, and lesser understood features from financial time series.

2.1 Fat tails and volatility

One of the most basic, and curious, empirical features of financial data is that the unconditional distribution of asset returns series are not Gaussian. [4] This is true at relatively high frequencies from minutes to months. At progressively longer horizons, return distributions slowly converge to Gaussian. For example, annual log differences of stock prices in the United States are very close to Gaussian. The speed of convergence to Gaussian distributions may differ across asset classes, but this is a stylized pattern for many series.

Much has been made about exactly what is going on in the tails of these distributions. Some of this has been driven by practical concerns coming from both risk management and derivative valuation, along with technological tools from extreme value theory. Other interest has come from the physical sciences where data descriptions based on "power laws" is relatively common. Researchers have attempted to uncover the shape in the distribution tail, approximating it with a form that looks like,

$$Prob(r > x) = Ax^{-\alpha}, \tag{2}$$

where α is known as the shape parameter. (LeBaron & Samanta 2005) present some recent evidence on α and references to the large literature on estimation, and some of the difficulties in precisely determining α. For daily return data, a value near $\alpha = 3$ is common. As one moves to longer horizons, estimated values of α steadily increase. The speed of convergence in the return distribution as the time horizon increases is an interesting and under explored area of finance.

After 40 years the causes for fat tails in return distributions are still not well understood. One would hope for some version of the central limit theorem to aggregate idiosyncratic noise across either people or time yielding Gaussian returns. Obviously, some form of dependence is keeping this from happening. This dependence could appear in correlation across behavior and trading strategies. Exploring this dependence is one of the crucial goals of the agent-based models which will be discussed later.

A second feature of returns data is the persistence of volatility which is interconnected to fat tailed returns. Most all financial series have persistent variances which are easily seen by looking at the autocorrelations of either the absolute or squared returns. These are positive and are significant for many lags. These changing

[4] This result goes back to the 1960's and (Mandelbrot 1963) and (Fama 1965).

conditional variances lead to models of the form,

$$r_t = \sigma_t e_t \tag{3}$$

where e_t is independent, identically distributed noise, and σ_t^2 is some kind of changing conditional variance. The large family of ARCH/GARCH models is an example of this. It is well known that models in this class will possess fat tails, but they are only an empirical description, and not a true behavioral mechanism explaining the existence of these distributions. Also, the distributions of estimated residuals of these models, e_t, are often fat tailed themselves, suggesting that changing variances alone do not give the whole story.

2.2 Extreme persistence

Volatility is not simply persistent. It is extremely persistent. Autocorrelations in absolute returns can be significant out at time lags of nearly two years. There is growing evidence that volatility is probably a long memory process possessing a slow hyperbolic decay in its autocorrelation pattern.[5] This deepens the puzzle as to what is causing this in financial data. One needs a persistence mechanism that will drop off sufficiently slowly to generate these types of patterns. As mentioned previously, this may be an indication that dynamics on multiple length scales are at work in financial markets. Trading volume is contemporaneously correlated with volatility, and also appears to be a good candidate for a long memory process as well.

Beyond volatility and trading volume, there are several other financial series which appear to be possible long memory candidates. These include dividend/price ratios, interest rates, and interest rate differentials. These again indicate that in all cases something may be happening on multiple time scales in terms of trader behavior, or the series may be subject to some nonstationarities or policy/regime changes. In all cases these features may make forecasting tricky.

2.3 High frequency/microstructure

The past decade has given rise to the field of empirical microstructure. This field looks at the dynamics of financial markets at very high frequency, often recording every trade. Patterns such as fat tails and persistent volatility are present in the high frequency data too. The data also reveal interesting daily seasonal patterns such as increased trading volume at the beginning and end of the day. In terms of dynamics, high frequency data also display some curious features. If one classifies trades as buyer initiated or seller initiated depending

[5] (Andersen, Bollerslev, Diebold & Labys 2003) is a good example of this work.

on how the order was executed in an electronic trading system, you will see extreme persistence in the type of trades. In other words, there are long stretches of buyer initiated, and long stretches of seller initiated trades.[6] What is also interesting, and puzzling, is that these long stretches of one sided markets don't destroy the martingale properties of returns themselves.

Order flow in foreign exchange markets has also been recently scrutinized. In several different papers Evans and Lyons use a similar kind of measure to show that order flow can be used to forecast exchange rates, and also macro economic variables.[7] These results are still early, and it will be interesting to see how well they hold up to replication on other series and time periods.

Higher frequency order flow data presents another interesting time series problem. Orders come in asymmetrically and are not evenly spaced. Their own time spacing may in itself be interesting and important. Exactly how to deal with and use this information on the dynamics of the trading process is an important and open problem.

3 Agent time scales

All of these results describe empirical properties of financial time series which are still not well understood. We have a large amount of econometric power dedicated to quantifying and understanding these phenomenon, but will still do not have many explanations for exactly where these features are coming from. A promising route to finding models that generate stock returns with real world properties is to use heterogeneous agent-based financial markets.[8]

The objective in building these models is to understand the dynamics of a heterogeneous set of boundedly rational agents trading securities in a simulated financial market. These agents should be close in spirit to the concept of building econometricians who are working to understand the world around them.[9] If this is the case, we should also accept the fact that they should not have solved some of the difficult problems that we observe in the real data. In most of my early examples of this I have assumed that the crucial problem they have not completely solved is that of whether the series they are looking at are stationary, and therefore what is the appropriate amount of data to use from the past when making decisions.

Examples and details on this type of model are contained in (LeBaron 2002). In this short note, I will

[6] These results are summarized in (Lillo & Farmer 2004).

[7] An example of their work on this is (Evans & Lyons 2004).

[8] See the website of Leigh Tesfatsion at Iowa State for an excellent summary on agent based models in economics, http://www.econ.iastate.edu/tesfatsi/ace.htm. Also, agent-based models in finance are described in (LeBaron forthcoming 2006).

[9] This is in the spirit of (Sargent 1993).

only comment on how the design and experiments are closely aligned with the problems of time scales. The market is composed of a group of agents that are making portfolio decisions between a single risky asset and a risk free asset. Their portfolio weights can depend on many pieces of past information, and they will chose dynamic portfolio strategies from a set of strategies. The key question in this process is "how should they make this decision?" Since the market is a complex heterogeneous set of agents there is no simple deductive scheme for deciding on an optimal trading strategy, so this becomes a purely empirical question. One might think that then it is clear that agents should use as much data as possible from the past to make decisions, but this imposes beliefs about stationarity on the agents, which given the evidence from the real world might be difficult to believe. I take a less constrained approach by having competing conjectures about appropriate amounts of useful past information coexisting in these markets. There will be agents who assume they need 50 years of data trading alongside agents who might only be using the last 6 months of simulated data. The market test selects across these agents in terms of the distribution of wealth. If the 6 month agent's performance is inferior to the others they will decrease to zero wealth, and disappear from the market.

Two major results have emerged from this model. First, agents with very different time scales can survive in the market together. The shorter horizon agents are not driven out of the market. Second, the dynamics of the interacting types yields many of the previously described empirical features. Among these are fat tailed return distributions, and long memory persistence in volatility and volume. Further experiments have shown that it is necessary to have agents with many time scales to get these results.

4 New econometric approaches for the future

The world of finance is fortunate to have data sets that span a wide range of frequencies and time scales. Using this range to its fullest extent will be necessary to estimate and validate the class of models just mentioned. They contain large numbers parameters which could give them many degrees of freedom in traditional testing settings. However, at the moment they may be the most likely candidates to generate features which occur on many times scales. This is natural since they are built from agent/decision maker populations which are somewhat confused about the appropriate time scale from which to analyze their data.

Techniques will need to be developed to test large models which are evolutionary in nature and do not live at any one special frequency. It will drive some currently separate fields closer together. For example, market microstructure work will have to become more closely aligned with more long range macro finance work. It still remains an interesting and open question whether there are features that spill over from the microstructure world to long horizons. The goal of understanding what drives the unusual features from financial time series is a big one. However, solving this is essential to understanding the big questions from finance. How and whether markets stay close to the ideal of market efficiency, and are there better measures of risk that can be accurately quantified in a large range of markets?

References

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **96**, 579–625.

Dacorogna, M. M., Gencay, R., Muller, U. A., Olsen, R. B. & Pictet, O. V. (2001), *An Introduction to High-Frequency Finance*, Academic Press, San Diego, CA.

Evans, M. D. D. & Lyons, R. K. (2004), Exchange rate fundamentals and order flows, Technical report, University of California - Berkeley.

Fama, E. F. (1965), 'The behavior of stock market prices', *Journal of Business* **38**, 34–105.

Ghysels, E., Santa-Clara, P. & Valkanov, R. (forthcoming 2006), 'Predicting volatility:How to get most out of returns data sampled at different frequencies', *Journal of Econometrics* .

Granger, C. W. (1980), 'Long memory relationships and the aggregation of dynamic models', *Journal of Econometrics* **14**, 227–238.

LeBaron, B. (2001), 'Stochastic volatility as a simple generator of apparent financial power laws and long memory', *Quantitative Finance* **1**, 621–631.

LeBaron, B. (2002), Calibrating an agent-based financial market, Technical report, International Business School, Brandeis University, Waltham, MA.

LeBaron, B. (forthcoming 2006), Agent-based computational finance, *in* K. L. Judd & L. Tesfatsion, eds, 'Handbook of Computational Economics', Elsevier.

LeBaron, B. & Samanta, R. (2005), Extreme value theory and fat tails in equity markets, Technical report, International Business School, Brandeis University, Waltham, MA.

Lillo, F. & Farmer, J. D. (2004), 'The long memory of the efficient market', *Studies in Nonlinear Dynamics and Econometrics* **8**(3).

Mandelbrot, B. B. (1963), 'The variation of certain speculative prices', *Journal of Business* **36**, 394–419.

Roll, R. (1987), 'R^2', *Journal of Finance* **43**, 541–566.

Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.

Tinbergen, J. (1933), 'The notions of horizon and expectancy in dynamic economics', *Econometrica* **1**(3), 247–264.

Tinbergen, J. (1939), 'The dynamics of share-price formation', *Review of Economics and Statistics* **21**, 153–160.