

BLACK sheep and WHITE



AUTHORS

Terje Lensberg
Professor

NHH – Norwegian School of Economics



Klaus Reiner Schenk-Hoppé
Marie Curie Visiting Professor

NHH – Norwegian School of Economics
Centenary Chair in Financial Mathematics
University of Leeds, UK



On September 15, 2008, Lehman Brothers, a large US investment bank, filed for bankruptcy protection. The event came to mark the onset of the financial crisis which recently spread to European sovereign debt. Since then, regulators around the world have implemented various emergency measures in an attempt to calm financial markets. Politicians became self-declared white knights, and short sellers were found to be the blackest sheep in the family of financial practitioners. Short sale bans and financial transaction taxes `to make speculators pay for their deeds' have been peddled as solutions to contain the crisis.

knights

Only five years before the turmoil in financial markets started, the Nobel Prize Laureate Robert Lucas declared «...macroeconomics in this original sense has succeeded: Its central problem of depression-prevention has been solved, for all practical purposes...». With the gift of hindsight, nothing could have been further from the truth. Indeed, economics experiences its own crisis as it contemplates alternatives to the standard rationality paradigm in modelling economic phenomena.

Evolutionary finance

Evolutionary finance applies Darwin's principle of natural selection to study trading behaviour and asset prices in financial markets. In this perspective, a financial market can be seen as a selection mechanism which transfers wealth to traders who are well adapted to the environment from traders who are less well adapted. The trading strategies of wealthy traders determine the prices of financial assets, first, because those strategies are backed by more wealth, and second, because the strategies of wealthy traders are more likely to be imitated by other traders. Wealth in financial markets is therefore the counterpart to fitness in biological systems.

Natural selection in financial markets is known to produce rational behaviour in the aggregate. The rationality assumptions that form the basis of the standard theory of financial markets can be seen as a proxy for this outcome. In many problems of interest, this proxy is excellent, in others it is not. One aim of evolutionary finance is to provide answers to questions on which the standard theory is silent, for instance how markets become efficient. Another aim is to study issues that cannot be analysed with the tools currently available within the standard approach. The recent financial crisis has raised a number of such questions.

Financial regulation

Since the outbreak of the crisis, financial regulators of many countries have

considered various measures to limit the negative effects of the crisis. The two most popular ones have been short-sale bans and Tobin taxes.

Short-selling means selling a stock that one does not own, and the popular opinion holds that short-sellers amplify market crashes by increasing the downward pressure on stock prices. Accordingly, a few days after the Lehman collapse, more than 20 countries introduced a temporary ban on the short-selling of financial stocks «to protect the integrity and quality of the securities market and strengthen investor confidence», as the US Securities and Exchange Commission (SEC) put it. In Europe, the temporary ban was reintroduced by Germany in May 2010, and by France, Belgium, Italy and Spain in August 2011, in attempts to stem the accelerating sovereign debt crisis.

A Tobin tax is a tax on financial transactions which is intended to curb speculation by making it more expensive to trade in financial markets. This tax has been used from time to time in the past, and has received fresh attention from regulators during the present crisis. Both the Nordic

and the European Councils have called for the introduction of such a tax, and in August 2011, the German Chancellor Angela Merkel and the French President Nicolas Sarkozy said they would make such a tax a priority for Europe.

One trouble with both these regulatory measures is that their effects are neither well documented nor well understood. This is mainly due to a lack of adequate theoretical models and a shortage of historical data. After the US short-sale ban was lifted in October 2008, SEC Chairman Christopher Cox told Reuters that «While the actual effects of this temporary action will not be fully understood for many more months, if not years, knowing what we know now, I believe on balance the commission would not do it again. [...] The costs appear to outweigh the benefits».

Computational approaches to the fore

Computational evolutionary finance employs techniques from computer science to carry out controlled experiments, for instance on the impact of market design or regulation. This is useful for regulatory bodies who would like to know more about

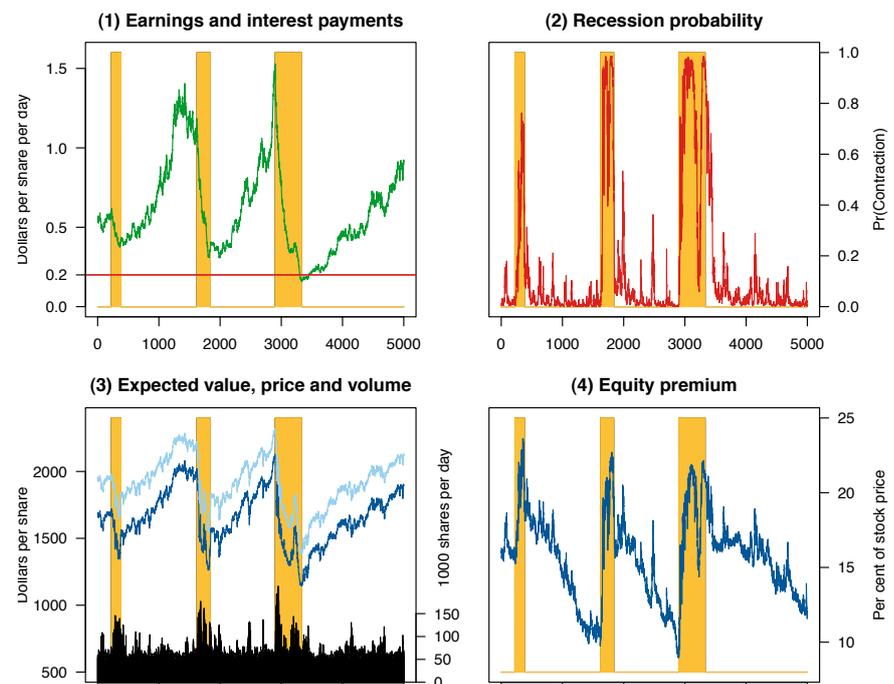


Figure 1 Simulation results of an order book market for equity in a firm with random earnings and unobservable business cycle regime.

the effects of new regulatory measures before implementing them in the market place. The typical model in this field consists of a detailed description of the market microstructure, and a large number of individual economic agents who make portfolio and trading decisions within that microstructure. The agents' trading strategies are represented as computer programs, and the model is solved by subjecting these programs to natural selection until the aggregate price process converges to a stationary process. Data from the model can then be analysed and compared with data from other experimental treatments with different market microstructure.

A model along these lines is illustrated in Figure 1. It contains two main parts: A financial sector and a real economy. The financial sector consists of 10,000 traders, who can invest in stocks and bonds issued by the firms in the real economy. They do this by submitting orders to a central stock exchange which maintains an order book and executes matching orders. The real economy is represented as one aggregate firm, whose main feature is an earnings process which is subject to short-run fluctuations and medium-run business cycles. It is modelled as an Ornstein-Uhlenbeck process with a Markov-switching attractor that can take on two values: A high value, representing a boom, and a low one representing a recession. The earnings process is illustrated in Panel (1) of Figure 1, where the yellow bars represent recessions and the white areas booms. Debtholders receive fixed interest payments, represented by the red, horizontal line, and stockholders receive the difference between earnings (green curve) and interest payments.

The traders can observe the earnings process, but not the state of the Markov process governing the business cycle. However, they do have access to a Bayesian estimate of the probability that the economy is in a recession. This probability is illustrated in Panel (2) of Figure 1.

In order to make money, the traders must use their information about current earnings, the Bayesian state probability, the order book and their own portfolio to identify profitable investment or trading opportunities. Over time, the trading strategies of

poor traders are replaced by copies or genetic recombinations of the trading strategies of more wealthy traders. After 3 million trading days, the gains from changing trading strategies are more or less exhausted. Then the simulation is terminated, the population is saved, and the model is restarted to generate data. Panel (3) of Figure 1 depicts the market price of one share (dark blue curve) along with the expected present value of the future cash flow per share (light blue curve). Traders in the model turn out to be risk averse: the stock trades at a discount relative to the expected present value of the cash flow it generates. This discount, shown in Panel (4) of Figure 1, is called the equity premium. It is higher in recessions than in booms, meaning that only investors with a large appetite for risk buy stocks 'when the cannons roar'.

Experimental results

Controlled experiments use simulations from a benchmark case with no short-sale ban and no Tobin tax, a case with a short-sale ban in place, and a case with only a Tobin tax. Each specification of the model is run for 100 different sample paths of the earnings process, corresponding to 100 scenarios of the real economy. Data are collected for each of the three experiments over 10,000 trading days for each of the 100 paths for the earnings process. Unlike in empirical research, we can re-run history with the benefit of creating a unique dataset of matched observations. This allows for comparisons of different regulatory regimes for an identical path of the real economy – the ultimate comparative analysis.

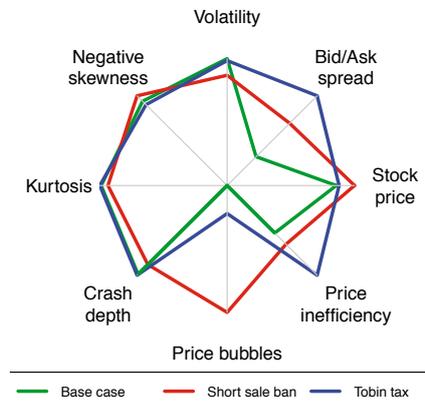


Figure 2 Indicators for the costs and benefits of different financial regulations of order book markets.

Figure 2 contains an overview of the results.

The radar plot provides information on the impact of financial regulation on order book markets along eight dimensions, all of which are central to the debate on the costs and benefits of regulation. Efficiency of the market is measured by liquidity (bid/ask spread) and price discovery; the strength and characteristics of price fluctuations are quantified using volatility, negative skewness and kurtosis; and the dynamics of long swings over the business cycle is captured by price bubbles and the depth of market crashes.

Short-sale bans were imposed with the goal of reducing price fluctuations, especially the depth of market crashes. Our results confirm that a short-sale ban indeed does have that effect. But the model also highlights the drawbacks. A short-sale ban increases the frequency of bubbles and leads to a generally overvalued stock market. The efficiency of the price discovery process is reduced, and so is liquidity.

Proponents of a Tobin tax have argued that it will reduce speculation and improve financial stability. These views are not supported by the model: A Tobin tax has no effect on the depth of market crashes, but large negative effects on price discovery and liquidity. The consequences are less trade and less informative prices; neither effect is positive.

Although the news seems to be one that will be met with delight by professional investors (and annoyance by politicians), there is more in store using our computational evolutionary finance approach. The model allows studying other modes of taxation such as a levy on market orders but not on limit orders or the introduction of a progressive tax scheme with an annual deduction.

Computational issues

Access to high-performance computing resources is a prerequisite for carrying out this type of experiments. The model sketched above has two main features which consume computing time. The first one is the order book, which, to capture the usual price-time priority of orders, must be represented as a tree with a branch for each stock price, followed by a sub-tree consisting of the orders submitted at that

price. Since the price has priority over order submission time, the priority of orders at one price is independent of their priority at other prices. One can therefore speed up the order book maintenance by implementing it as an array of sub-trees, one for each price.

The second computationally costly feature is associated with the genetic programming algorithm (GP) which is used to evolve trading strategies and solve the model. There are many variants of GP algorithms, we use a steady-state algorithm with tournament selection, which works along the following lines:

1. **Initialization.** Randomly generate a population of 10,000 trading programs.
2. **Trading.** Traders are randomly and repeatedly called upon to submit orders which are handled by the stock exchange.
3. **Tournament.** Randomly select 4 programs and rank them by (discounted) wealth.
4. **Reproduction.** Replace the two poorest programs by copies of the two wealthiest ones.
5. **Crossover.** With some given probability, swap random subsets of program instructions between the two copied programs.
6. **Mutation.** With some given probability, replace a randomly selected instruction in some copied program by a randomly generated instruction.
7. Go to 2.

The algorithm has a few additional features that are worth mentioning: (a) In step 3, traders are selected by first choosing the size of a local neighbourhood of traders, and then randomly selecting traders for the tournament from that neighbourhood. This form of local competition is known to improve the ability of the algorithm to generate innovative and potentially superior behaviour; (b) In step 2, a machine code representation of the programs is used for fast computation of trading decisions; and (c) in steps 5 and 6, a fixed-length byte code representation is used to

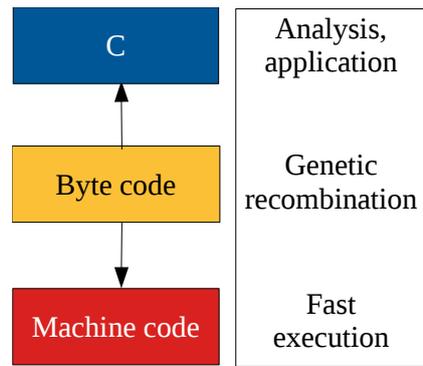


Figure 3 Program representation in GP-algorithm.

simplify the genetic recombinations. A fast built-in compiler translates byte code to machine code which is then typically used thousands of times before it needs to be recompiled. Our software also contains a byte code disassembler which generates C code that can be used for visual inspection of the evolved programs.

Parallel processing

Another feature of the GP model discussed so far is a tight interconnection between the individual agents of the population. It arises because the traders do business with each other on one common marketplace. This type of model therefore does not lend itself easily to parallelization due to the large amount of information that would have to be exchanged between the traders and the marketplace¹.

On the other hand, engineering applications of GP are well-suited for parallelization. That includes financial engineering applications such as risk management and computerized trading systems. Developing such systems typically involves solving some ill-structured optimization problem, which is a natural habitat for GP. To solve these types of problems, one can use a large number of autonomous sub-populations deployed on separate worker nodes, and let them exchange good candidate solutions with their neighbours from time to time. A master node can be added to collect information

about the progress of the worker nodes and perform other administrative tasks. Such a structure eliminates waiting time on the worker nodes and yields a high capacity utilization rate for the whole system.

Lessons learned

Our research applies the new field of computational evolutionary finance to study the effect of short sale bans and transaction taxes on financial market stability. The approach provides an unparalleled detailed model of order book markets and offers new insights into their dynamics. Neither financial transaction taxes, nor the emergency actions taken in 2008, and more recently in 2011, by imposing short sale bans, are capable of delivering what politicians desperately seek; calm and quiet markets. In the figurative nutshell: Short-sellers and active traders are not black sheep – no matter what self-declared white knights want to make believe.

Outlook

The approach combines natural selection with models of market interaction. But its potential for applications goes well beyond the specific model of financial markets presented here. Our current research looks at disparate issues such as the pricing and hedging of exotic options, and the impact of market fragmentation on the evolution of the market ecology and traders' investment skills. ●

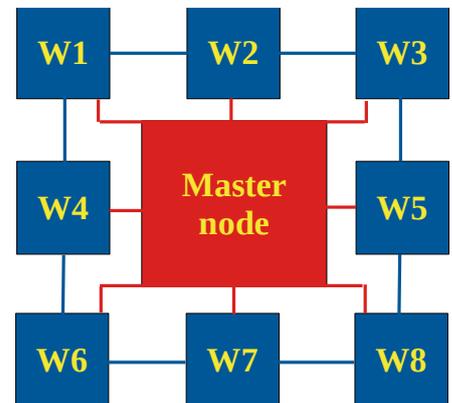


Figure 4 Structure of parallel GP-algorithm.

¹ In real markets, this information flow is routinely handled by exchanges and other market venues. For example, the Chicago Mercantile Exchange processes some 13 million contracts per day. By comparison, the exchange in our model handles an order flow of that magnitude in less than 20 seconds of wall clock time. Obviously, that is possible only because all communication takes place on the same chip.