Evolutionary simulation of hedging pressure in futures markets

J. Duke and C. D. Clack

***Abstract*— We present a real world application that models a financial futures market. The agent-based simulation includes speculator agents each of which uses a Genetic Algorithm to improve its profitability in the market. This is a realistic simulation whose rates-of-return distribution is similar to those of real futures markets such as corn and FTSE100 futures.**

The futures markets have never before been simulated to this level of detail, and the simulation is used to test the long-held belief that speculators are more profitable if they incorporate "hedging pressure" into their price calculations — essentially, the use of market knowledge about supply and demand. Surprisingly, we show that hedging pressure cannot be used to improve profits for speculators.

I. Introduction

Futures markets are a long-standing and essential part of the financial world. Futures trading volume has increased substantially as organisations become more sophisticated in their management of risk and as speculators try to cash in on the resulting supply and demand.

The price of a futures contract in the real world is rarely equal to the theoretical expected future spot price as given by the cost-of-carry model (see [11]). Prices can vary either side of this theoretically "correct" price according to current market sentiment; these variations are known as "risk premia". There is no consensus on the precise composition of these premia, but previous work (e.g. [12], [3]) has analysed historical data and found that systematic risk and hedging pressure (an excess of supply or demand caused by "hedgers" — underlying producers and consumers) seem to play a part.

In this paper, we explore hedging pressure using an evolutionary agent-based simulation of a futures market. This approach introduces the possibility of feedback between competing agents, and thus goes beyond any straightforward analysis of historical data. The simulation contains agents representing hedgers and speculators (see Section III), and tries to evolve accurate pricing rules — that is, pricing rules that can take account of any risk premiums present in prices — using a genetic algorithm. The simulation produces a returns distribution similar to real futures markets, and exhibits clear signs of price pressure caused by hedgers. However, examination of the price predictions made by various pricing models in the simulation shows that the models that try to account for this hedging pressure are no more accurate than the models that do not.

The authors are with the Department of Computer Science, University College London, Gower Street, London WC1E 6BT (email: clack@cs.ucl.ac.uk).

II. Previous work

A. Agent-based simulations

Understanding financial markets is a complex task. Trying to extract particular features of a market into a tractable mathematical model inevitably means losing other essential features, making it hard to build up a complete picture of market behaviour. To try to study market interactions and emergent behaviour, researchers use computer simulations that allow independent "agents" to trade with each other via an artificial marketplace. For review of agent-based work in finance and economics, see [15] or [1]. These simulations allow researchers to conduct controlled experiments, something that is impossible in real financial markets.

There are many of these artificial markets. One of the most studied is the Santa Fe Artificial Stock Market [2], [14]. This uses agents that make price predictions using a variety of rules, with each rule only applying in certain market conditions. Agents can choose between a risky asset and a risk-free asset, and evolve their rules using a genetic algorithm. Closely related to the Santa Fe Institute, BiosGroup Inc. developed an agent-based model of the NASDAQ market that including the use of reinforcement learning within agents [7].

Chiarella and Iori [6] create a very simple agent-based model that is designed to examine market microstructure. The results indicate that fundamental, technical and noise traders are all necessary to generate realistic market behaviour — this seems to back up the results in [10], and we include all three types of agent in our simulation.

Some attempts have been made to simulate derivatives markets. Most notably, King et al [13] (building on [17]) attempt to simulate an options market. De la Maza and Yuret [8] specifically simulate a futures market, creating a model with futures contracts that always expire the following day. Agents maintain their own strategies, some of which are visible to other agents and some of which are not, and these strategies evolve over time using a genetic algorithm. The market here is very simplistic, and the primary interest is the heterogeneous nature of the agents, and specifically whether it is possible for some agents to be consistently profitable.

B. Hedging pressure in futures markets

Hedging pressure is an excess in supply or demand, causing futures prices to deviate from expected spot prices. Traders expect to earn (or pay) risk premia according to the excess [9], but the composition of these premia is unknown.

Bessembinder [3] finds evidence that hedging pressure is a determinant of premia in some futures markets. [9] use a similar model and suggest that risk premia depend on hedging pressure not only from the underlying market, but also from other related futures markets.

These studies rely on straightforward mathematical analysis of historical futures returns. The results suggest that taking hedging pressure into account on a day-to-day basis could lead to better price predictions and hence greater profit for speculators. However, exploring the effects of hedging pressure in day-to-day trading is not possible by analysing historical data: the only available data detailing the proportion of hedgers in futures markets are of weekly granularity, and only cover U.S. markets[1]. In addition, if a large proportion of speculators were to make predictions, and thus trades, while accounting for hedging pressure, it is possible that any hedging pressure effects would be reduced, or disappear altogether.

In order to explore hedging pressure in more detail we have therefore attempted to create an agent-based simulation of a futures market, containing agents representing both hedgers and speculators. The simulation will allow us to examine the effect of hedging pressure on daily price predictions, made without hindsight, and also to see if price feedback from interactions between traders would negate any advantage that accounting for hedging pressure may yield. This information would be of help to the many speculators who make regular updates to their positions and are constantly engaged in an "arms race" to gain more knowledge of the market than their rivals. These speculators often use models that build up a view of the market from many different factors, and if hedging pressure could be shown to be a useful addition to any such model, then even weekly updates on hedger numbers would be of value.

III. Futures contracts

Futures are an example of a financial instrument known as a *derivative*, whose value is derived from an underlying asset. A *futures contract* (or "future") is a contract, traded on a futures exchange, to buy or sell that underlying asset at a certain date in the future, at an agreed price[2].

Futures originated as a mechanism for producers to lock down a price for the future sale of goods that they had not yet produced; and likewise for consumers to lock down a price for goods that they intended to buy in the future. So, for example, a wheat farmer can decide how much to plant without worrying that a shift in wheat prices will leave him out of pocket or unable to sell; and a miller can guarantee a supply of wheat and remove worries that changing prices will leave him unable to afford the quantity he requires.

A future can therefore be used to reduce *risk* for one or both parties in the contract; a process known as *hedging*. Hedging creates demand to buy and sell futures contracts, which allows traders known as *speculators* to take positions in the futures market with no intention of ever buying or selling the underlying asset—they hope to profit by taking on the risk offloaded by the hedgers.

[1]The Commodity Futures Trading Commission's weekly Commitment of Traders reports.

[2]Futures are specifically those contracts traded on a futures exchange; any other agreement between two parties to buy or sell at some point in the future is called a *forward* contract.

A. Pricing futures: the cost-of-carry model

The simplest model of futures pricing is called the *cost-of-carry model* (see [11]). This model states that a future should be priced so as to prevent profitable arbitrage opportunities. For example, it would be possible for a trader to buy the underlying asset now, and also take out a futures contract agreeing to sell that asset at a particular price at some point in the future. The price on the futures contract should therefore reflect the cost of buying the asset and and storing it until the expiry of the future, else a trader could lock in a guaranteed profit. (Agreeing to sell the underlying asset is known as taking a *short* position; an agreement to buy the underlying asset is known as a *long* position.)

B. Futures in the real world

1) Risk premiums and futures speculation: In the real world, futures prices do not follow the theoretical prices of the cost-of-carry model. The time gap between the buying or selling of a contract and the expiry of that contract introduces uncertainty. This uncertainty means that futures prices include so-called *risk premiums*, which depend on the current market sentiment. For example, if the market is unsure about the future availability of the underlying asset, there will be more demand for long futures contracts as manufacturers try to guarantee their future supply. This demand drives up futures prices. Speculators try to anticipate this market sentiment, taking out long futures positions in the hope of profiting from the rising prices.

Speculators use many different strategies to try to predict price movements: some rely on mathematical (or *technical*) analysis of price series; some look at real-world data such as levels of employment or interest rates (known as *fundamentals*); many use a combination of the two.

2) Types of traders in real markets: In most real futures markets, speculators outnumber hedgers. The Commodity Futures Trading Commission publishes weekly surveys of the nature of every trader holding a position in any U.S. futures market. Traders are classified as commercial (i.e. hedgers — traders with real interest in buying or selling the underlying asset) or non-commercial (i.e. speculators). For an in-depth study of the way futures traders behave, see [18].

3) Futures trading: In a real futures market, contracts with various different expiry dates will be available for trading. However, at any particular time, most of the trading volume will be in only one of those contracts; usually, the contract that is closest to expiry. As the expiry date of that contract approaches, speculators and other traders who do not wish to hold their positions until expiry will get out of those positions, and take up similar positions in the next-nearest contract. This movement from one contract to another is known as "rolling".

4) Back-adjustment: For technical pricing models, traders need a continuous historical price series, without needing to account for the fact that the prices are for different futures contracts and thus not directly comparable. They therefore transform the prices by a process called *back-adjustment*.

Going backwards in time, at the point of "rolling over" the price difference between the old and new contract is calculated. This adjustment is then added to all the prices for the previous contract. This process then continues with the adjustment being accumulated with each contract, i.e. the difference between the two contracts on the next roll day is added to the current adjustment, and then that is applied to all prices from that day backwards, and so on, to the beginning of the price series.

5) Orders: A real market usually allows many different types of order to be placed. The most commonly used orders are called limit orders and market orders.

A limit order is an order where a trader can specify a "limit" price, and know that the order will only be executed if the execution price is at least as good as that limit price. Limit orders also have a "size", specifying the number of contracts offered or desired at that price. Limit orders to buy are called *bids*: limit orders to sell are known as *asks*.

Limit orders are stored by a mechanism known as an *order book*. The order book keeps track of all limit orders placed, ordered by how competitive the prices are. This puts the highest bids and the lowest asks at the "top" of the order book. The highest bid at any time is known as the "current bid", and the lowest ask is the "current ask". A limit order will expire, i.e. be removed from the order book, after a set length of time if it has not been executed.

A market order is an order to be executed immediately, at whatever price is currently available; in other words, the order will be matched against the limit orders currently stored in the order book. If a trader places a sell (or buy) market order of a size greater than the size of the current bid (or ask), then the order book will continue matching against the bids (or asks) until the market order has been filled or until there are no more bids (or asks) remaining.

IV. Hedging pressure in an agent-based simulation

A. The simulation

The simulated futures market contains agents representing both hedgers and speculators (as in [4]), as well as "noise" agents that represent traders with no firm strategy and that help to provide liquidity[3]. There is no intent to simulate any specific real-world market. Rather, the simulation is merely a generic futures market; the actual underlying asset is unimportant for the purpose of this work.

1) Timesteps: The simulation runs over a series of discrete timesteps; each timestep can be thought of as representing one "day". To avoid having to simulate a market in the underlying asset, the spot price (i.e. the price of the underlying asset for immediate delivery) begins at a particular value and is adjusted at each timestep by adding small random quantity drawn from a normal distribution with low variance.

[3][10] and [6] also suggest that such agents help to generate realistic market behaviour.

2) Hedgers: Hedger agents are either *long* or *short*, and only take positions on that side of the market. They place orders at random intervals, with some hedgers trading more frequently than others, and buy or sell a fixed quantity in each trade. They always hold their positions to the expiry of a contract, because they are in the market with the intention of making or taking delivery of the underlying.

3) Speculators and pricing models: Speculator agents make price predictions using a weighted average of the prices predicted by four simple pricing models, and make trades based on those predictions. The four models are: a simple technical strategy (linear extrapolation); the standard cost-of-carry model; the technical strategy with an adjustment made to account for the number of hedgers in the market; and the cost-of-carry model with the same adjustment. The adjustment assumes that when long hedgers are in the majority, the price will be driven higher, while a majority of short hedgers will push the price lower; in other words, it attempts to account for hedging pressure.

The total number of hedgers in the market is constant, but the number of long and short hedgers varies throughout the simulation, with either long or short hedgers always in the majority. A speculator always has up-to-date knowledge of these numbers, but not of trades made or positions held.

4) Market microstructure and futures contracts: Agents can place limit and market orders. Limit orders are stored in an order book. Each agent has a level above (below) their predictions at which they are willing to place market orders to sell (buy). If the current bid (ask) is below (above) this price, then the agent will instead place a limit order at that price. This methodology is taken from [6].

5) Rules and Evolution:

Agent rules: Each speculator agent forms price predictions by using a set of rules that match market conditions to combinations of their four pricing models. These rules are loosely based on those used in the Santa Fe Artificial Stock Market [2]. The state of the market at any time is characterised using a set of seven true-or-false statements:

1) Current futures price $>$ 5-day moving average
2) Current futures price $>$ 10-day moving average
3) Current futures price $>$ 100-day moving average
4) Current futures price $>$ 500-day moving average
5) 5-day moving average futures price $>$ 10-day moving average
6) 10-day moving average futures price $>$ 100-day moving average
7) Current futures price $>$ expected future spot price according to cost-of-carry model

At each timestep, a seven-bit binary string is generated summarising the current market state, with a `1` for a true condition and `0` for a false condition. For example, the string `1000000` would mean the current price must be above the 5-day average but below all other longer-term averages, suggesting a sharp recent increase in the futures price.

Some combinations of bits can never occur. For example, if the current price is below the 5-day average but above the 10-day average, then the 5-day average must be above the 10-day average. After taking all such impossibilities into account, there remain 72 possible market states.

To make predictions, speculators search through their ruleset looking for rules that match the current market conditions. In every rule, each market characterisation statement is associated with a `1`, a `0` or a `#`: a `1` means that the statement must be true for the rule to match, `0` means that the statement must be false for the rule to match, and `#` means that the statement is ignored (i.e. a "don't-care" symbol). For example, the market condition string `1000000` would be matched by the rules `100#00#`, `#000000`, and `1#0#0##`, but not by `100010#` or `00#0000`. Each agent will be required to include an all-`#` rule that matches regardless of current market characteristics, so that an agent always has at least one rule with which it can make a prediction.

A rule then associates particular market conditions with a set of four weights. Each weight corresponds to a particular pricing model. Also associated with each rule is a measure of the accuracy of that rule, tracking whether the rule has historically made accurate price predictions.

At the start of the simulation, all values except accuracy will be initialised randomly; accuracy will be set to 0. The four weights will each initially lie in the range $[0, 1]$, and then will be normalised to sum to 1.

Overall, a speculator agent consists of fifty of these rules, which should be sufficient to cover a good range of market states and model weightings without being too computationally intensive, and the all-`#` rule. Each rule has seven market-characterising bits, four floating point numbers representing the weights associated with each pricing model, and a floating point number for the accuracy of the rule.

Evolution: A genetic algorithm is used by the speculators to try to improve their rules over time by discarding historically inaccurate rules and generating replacements from combinations of accurate ones. The evolutionary process is similar to that used in the Santa Fe Artifical Stock Market [2].

As the simulation runs, the accuracy of the price predictions is monitored. After a sufficient period of time, it is possible to identify those rules that are consistently inaccurate, and those rules that never get used. Agents evolve by removing such rules and trying new rules in their place, using the process of selection, recombination and mutation.

Selection happens by first creating a pool of every rule used by every agent, and ranking them according to accuracy[4]. The top 20% of this entire set of rules, i.e. the most accurate global 20% of rules, are then used as a pool from which *parents* are selected.

Each agent then evolves independently. An agent will first discard the worst 20% of its rules. Each of these rules is replaced by a new rule, formed by recombining of a pair of rules from the global pool of parents. Recombination happens as follows: first, a new bit string will generated using uniform crossover on the parent bit strings[5]; next, the floating point numbers representing the weights will be crossed over using an accuracy-weighted average of corresponding values in the two parents; and finally, the accuracy value of the child will be the average of the accuracy of the parents.

Finally, the agent's rules may be mutated, i.e. changed slightly at random. For the bit string, each bit will be changed with a probability of 0.02; any bit to be changed has an equal chance of changing to either of the two values to which it is not currently set[6]. The weights are mutated by "creep" mutation with a probability of 0.02. Any weight to be mutated is changed by adding a number drawn from a normal distribution with variance 0.02. The weights are renormalised to 1 after mutation.

If models that account for hedging pressure do indeed lead to a sustained improvement in prediction accuracy, then rules that give greater weighting to these models would be expected to emerge as dominant in the population. In the real world, traders learn from their mistakes and change their trading methodology over time. A pricing model that yields consistently poor results is unlikely to be used for very long. In analagous fashion, speculator agents in the simulation use evolutionary methods to try to adapt their pricing models to produce more accurate predictions.

6) Main loop: Each timestep of the simulation proceeds as follows:

1) Update the accuracy of rules that matched the market conditions in the previous timestep, by comparing their predictions with the actual price generated.
2) Get market conditions for this timestep.
3) Go through all the agents in random order and ask each one to make trades given the current market conditions, price history and order book.
4) Get the latest price quote from the order book and add to the historical price history.
5) Randomly decide whether to switch between a majority of long hedgers and a majority of short hedgers.
6) If necessary, tell the order book to introduce a new contract, and apply back-adjustment to historical prices.
7) If necessary, go through the evolutionary process adjusting the rulesets of the speculator agents.
8) Add a small random variation to the spot price.
9) Proceed to the next timestep.

B. Measurements

1) Pricing model weights: Agents make price predictions by taking a weighted average of the prices predicted by their four models. If the model combinations were random, the average weight of each of the four models across all agents would be $\frac{1}{4}$. However, the evolutionary process will favour models that produce accurate predictions at the expense of

[4] A rule corresponding to market conditions that have never occurred is deemed to be of lower accuracy than a rule that has matched just once, no matter how inaccurately. This ensures that rules that are never used get eliminated. The warmup period at the start of a simulation makes sure that a good variety of market conditions are generated before any evolution happens, giving rules every chance to be tested.

[5] "Uniform" crossover means that each bit in the new bit string will be a corresponding bit from one of the two parents, chosen at random with equal probability.

[6] For example, if a bit is currently `1`, it could change to either `0` or `#`.

less accurate ones, so it would be expected that any more accurate model would, over time, emerge with a higher-than-average weight. To monitor this, the weights used to make predictions at each timestep are collected for later analysis. It would be reasonable to expect a pricing model to be more successful under some market conditions than under others. Information about the market conditions at each timestep will therefore be collected alongside the weights used.

2) Processing of results: The results will be collected from simulations run for 28,000 timesteps — see Section V-A. There will be a warmup period to give the simulation time to generate a wide variety of market conditions, testing as many rules as possible before the evolutionary process begins. After that, the main part of the simulation will commence, with evolution happening at regular intervals. Finally, the results will be collected in a "warmdown" period of each simulation run, during which evolution will halt, keeping the set of trading rules fixed in each agent; at each timestep during this period, every rule that a speculator uses to make a prediction will be stored, along with the market conditions for that timestep.

The processing of these results for later analysis will then proceed as shown in Figure 1. The rules will be examined to find the ten most accurate speculator agents, i.e. the ten speculators whose rules were most accurate on average. The complete set of rules used by these agents will then be collected together. This set of rules will be separated according to the market conditions at the timestep when each rule was used, so that all the rules used when the market was in a particular state are together.

Now we take the set of rules for a particular market condition. Each rule contains four numbers which are the weights given to each pricing model. For each pricing model, the set of weights for that model from all rules in the set is averaged, to give us a final set of four numbers, one for each model. This is done for every market condition, resulting in four numbers for each of the 72 market conditions. For any one experiment, 50 such simulations will be run, to give a final total of 288 (72 market conditions $\times$ 4 pricing models) 50-element vectors.

V. Properties of the evolutionary simulation

A. Parameter choice and sensitivity

The warmup period had to test as many market conditions as possible, and after some testing, 4,500 timesteps was found to test over 70 of the possible 72 conditions on average, and this was judged to be good enough. The main part of the simulation was chosen to be 16,000 timesteps, with evolution occurring every 2,000 timesteps—while this may not seem to produce many evolutionary phases, it was found that beyond this point, the population of rules began to lack diversity, making further evolution ineffective. The warmdown period was chosen to be 7,500 timesteps, which was easily sufficient to test the rules that had been produced. This gave a total of 28,000 timesteps per run of the simulation, roughly equivalent to 80 "years".

The total number of agents in the simulation was chosen to be 100. This produces a good variety of strategies, gives the market sufficient liquidity, and keeps the running time of the simulation at a practical level. It is comparable to the number of agents used in similar work, e.g. [2], [16]. Running the simulation with 150 agents produced almost identical results, and reducing the number to 50 had little effect.

B. Evolution

The evolutionary process plays a major role in the simulation. The evaluation of pricing models depends on the idea that inaccurate rules will get discarded and better ones will be found. It was decided that agents would evolve simultaneously: every 2,000 timesteps the simulation is interrupted and every agent updates (evolves) its ruleset. Changing this so that agents have a 1 in 2,000 chance of evolving at every timestep had a noticeable effect, but did not significantly change the overall conclusions. It is likely that any difference was due to a faster rate of change, as accurate rules will be capable of "reproducing" sooner, and thus will come to dominate the parent pool more rapidly.

It may be that the success of a particular model will vary over time. For example, as a model becomes dominant (i.e. has a much higher weighting than the other models), it will be trading mostly against itself, which it may find more difficult than trading against others and so its accuracy will decrease and allow other models to re-emerge.

The fitness of a rule is based on its historical accuracy. Each rule stores a number representing this accuracy; the smaller the number, the more accurate the rule. At each timestep, any rule that matched the market conditions at that timestep will update its accuracy by taking a weighted average of two numbers: the previous accuracy measurement; and the absolute difference between the rule's predicted price and the actual price from the simulation. After some experimentation, it was found that making one weight substantially higher or lower than the other did have a significant effect on the results. For the simulations described here, it was decided to make the weights equal. The length of time taken for a change in the quality of a rule's predictions to be reflected in the accuracy measure depends on these weights; this choice balances the conflicting desires to have accuracies as up-to-date as possible and yet not entirely ignore historical information.

The rules of each speculator are initialised randomly to avoid producing agents that are too similar to each other. There is therefore no guarantee that the space of possible rules will be well covered. However, testing the simulation with rules initialised systematically — ensuring good coverage across the rulespace — yielded no difference in results. This suggests that there are sufficient rules and agents for random initialisation to be effective.

The contracts in the simulation are available for 32 timesteps each with a small overlap between contracts, and can be thought of as "monthly". The simulation was also tested with much longer "biannual" contracts and this did make a

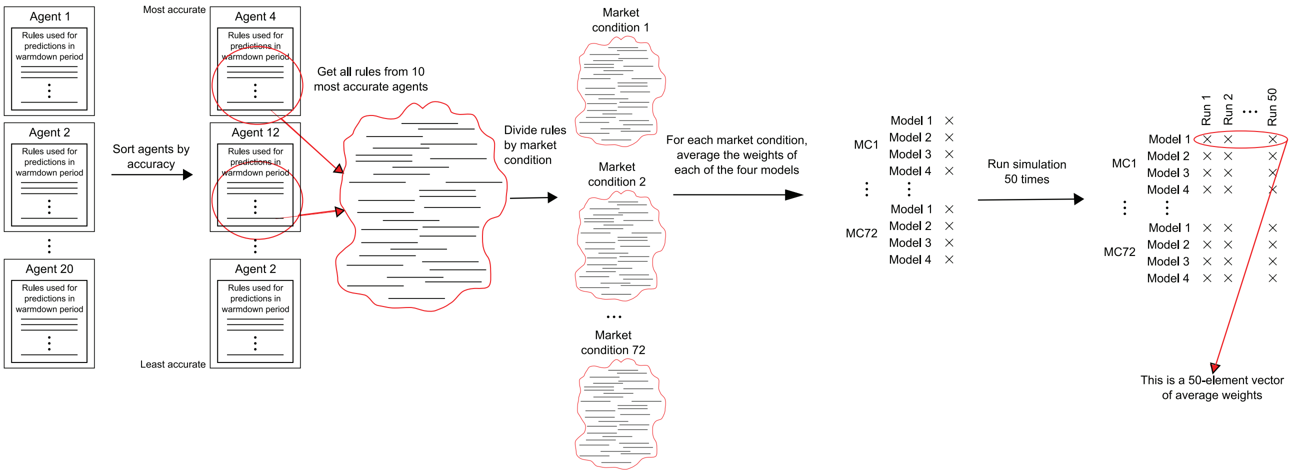

Fig. 1. Diagram showing the process by which the results of the simulation are broken down to a set of 288 50-element vectors for analysis. See Section IV-B.2 in the text for an accompanying description.

noticeable difference, but monthly contracts are a reasonable assumption for a futures market[7].

C. Stability and realism

The simulation must be both stable and realistic: in particular, the distribution of returns generated by the simulation should be similar to those found in real markets.

The returns distribution of a sample run of the simulation is shown in Figure 2, along with the returns distributions for various real futures markets. It is clear that the distributions of real markets vary widely. The simulation is probably closest to the distribution of corn futures. A ranked T-test comparison of 500 returns with the corn market gives a p-value of 0.44, indicating no significant difference in the two markets.

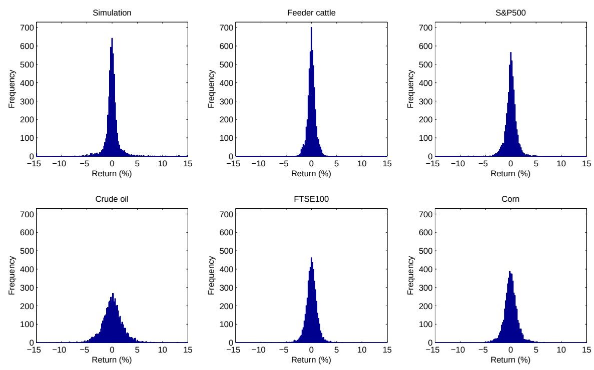

Fig. 2. The rates-of-return distributions for the simulation (top-left) and for various real futures markets. Each histogram shows 5000 daily percentage returns on identical axes, using identical bins. (Source: Commodity Research Bureau, www.crbtrader.com)

[7]Real futures markets with monthly contracts include crude oil, gold, copper, rubber, eurodollars, live cattle and the Mexican peso-U.S. dollar exchange rate.

VI. Hedgers in the Market

A. Hedgers on one side of the market

Having shown the simulation approximates some reasonable market, and having gained an understanding of how important the various parameters are and how the simulation behaves, we can change specific elements of the simulation in an attempt to tell us something about real markets.

The first test was to look at the basic effect of hedgers on prices. The simulation was first run with 5 long hedgers and no short hedgers, then with 5 short hedgers and no long hedgers. In both cases, the simulation ran for for 5,000 timesteps, without evolution, using 25 speculators and 75 noise traders. The resulting back-adjusted prices for the first case are shown in Figure 3. The graph shows that, in isolation, long hedgers cause an upward price drift. Similarly, short hedgers cause a downward price drift. This is exactly what would be expected according to the definition of hedging pressure.

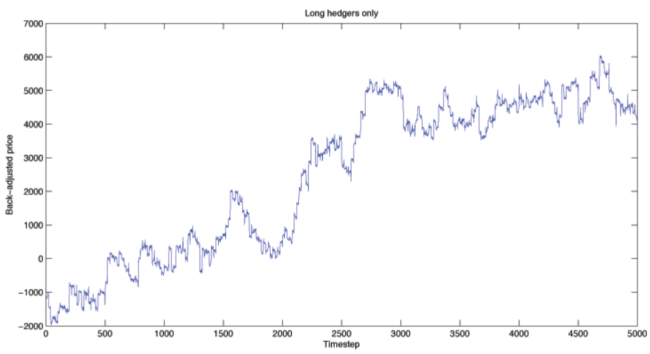

Fig. 3. Back-adjusted prices for a simulation with only speculators, noise traders and long hedgers. There is a clear overall trend upwards, as would be expected.

B. Equal numbers of hedgers on both sides of the market

The next test looked at the overall effect of hedger agents when there are equal numbers of hedgers on both sides of the market, minimising any possible hedging pressure effects. The simulation was run with a total of 10 hedgers, then with 20, 30, 40, and finally 50, always with equal and constant numbers of long and short hedgers. This removed hedging pressure adjustments from the price predictions. The remainder of the total of 100 agents was made up of 25% speculators and 75% noise traders.

For each number of hedgers, the simulation was run 50 times (28,000 timesteps each time). All settings were as described in Sections IV and V. During the warmdown period the simulation collected the average weights given to the four models (producing 50-element vectors as described in Section IV-B.2), and the predictions made by each model at each timestep. The average weights were then analysed to see how they were distributed between the pricing models.

For the purposes of this analysis, the market conditions were divided into three categories, chosen due to clear differences when looking at the results.

1) A strong recent downward trend.
2) No strong recent trend.
3) A strong recent upward trend.

With equal numbers of long and short hedgers, the weights given to the models that account for hedging pressure were roughly equivalent to the weights given to their counterparts that make no adjustment.

VII. Main Results

A. Hedging pressure

Finally, the effect of hedging pressure itself could be explored. The aim was to examine whether futures prices can more accurately be predicted when the proportion of hedgers present in the market is taken into account, and to see if price feedback from interactions between traders negates any advantage that may accrue.

The simulation was run many times using differing numbers of hedgers. If hedging pressure does have a significant effect on prices, then the models that try to account for this effect should be more accurate, and should thus evolve to have greater weight than models that do not take account of hedgers. It would also be expected that when the overall proportion of hedgers in the market is low, the models that take no account of hedgers will be more successful than when the number of hedgers is high.

This can be summed up with the following null hypothesis:

- The number of hedgers in the market has no effect on the relative success of the models in the simulation.

To test this, the simulation was run with 5, then with 10, then 15, 20, 25, 30, 35, 40, 45 and 50 hedgers, with the remainder of the total of 100 agents always made up of 25% speculators and 75% noise traders. The total number of hedgers will remain constant throughout any one simulation, but the number of long and short hedgers will vary randomly over time to allow for possible hedging pressure in both directions.

As in Section VI-A, the simulation was run 50 times for each number of hedgers, for 28,000 timesteps each time. During the warmdown period the simulation collected the average weights given to the four models (producing 50-element vectors as described in Section IV-B.2), and the predictions made by each model at each timestep. The average weights were then analysed to see which models were given most weight, and the predictions examined to see how accurate each model was and confirm that the most accurate models were indeed given most weight.

B. Results

Table I gives the average weights for each of the three categories (as in Section VI-B) for each possible number of hedgers. The weights show that the models that account for hedging pressure are consistently given *less* weight than their counterparts that make no attempt to account for hedging pressure. This is the opposite of the expected result.

Also apparent is that cost-of-carry models are always given greater weight than the pure technical models. However, they are given less weight as the number of hedgers goes up, particularly in market condition categories 1 and 3; the lost weight is spread between the two pure technical models. In market condition category 2, the difference between the models is less, and remains largely unchanged no matter how many hedgers are in the market.

VIII. Conclusions, discussion and further work

There are two issues here: the quality and efficacy of the futures market simulation, and the questions surrounding hedging pressure within that simulation. The overall structure of the evolutionary simulation seems to be adequate to create a vibrant marketplace displaying plenty of liquidity and allowing agents to use a number of different strategies.

The questions surrounding hedging pressure are harder to answer. The results in Section VII-B suggest that in a marketplace where some proportion of traders are trying to account for hedging pressure effects, those effects do not exist, and that the traders could probably produce more accurate predictions without making these adjustments.

Were the pricing models were too simple to exploit the hedging information? The simplicity and absolute accuracy of the pricing models is not what is being tested here: what matters is whether including an adjustment for possible hedging pressure effects produces more accurate predictions than when not including this adjustment. In the tests here, the adjustment did not lead to even slightly more accurate predictions, as might have been expected.

It is curious that an increased number of hedgers leads to lower overall accuracy, and better relative accuracy of the pure technical models in comparison to the cost-of-carry models. Perhaps the most likely explanation is that with more hedgers in the market there are fewer speculators, and thus fewer different rules. This would reduce the size of the

TABLE I

RESULTS FOR SIMULATIONS RUN WITH A VARYING NUMBER OF HEDGERS, ALWAYS WITH A MAJORITY ON ONE SIDE OF THE MARKET OR THE OTHER. THE TABLE SHOWS THE AVERAGE WEIGHTS GIVEN TO THE FOUR MODELS (PURE TECHNICAL, COST OF CARRY, PURE TECHNICAL + HEDGING PRESSURE, AND COST OF CARRY + HEDGING PRESSURE) IN PREDICTIONS MADE BY THE TEN MOST ACCURATE SPECULATORS DURING THE WARMDOWN PERIOD. THE MARKET CONDITION CATEGORIES ARE SPECIFIED IN SECTION VI-B.

	Market condition category 1			
No. hedgers	PT	CoC	PT+HP	CoC+HP
5	0.181	0.337	0.155	0.327
10	0.171	0.341	0.160	0.328
15	0.181	0.337	0.156	0.326
20	0.180	0.343	0.159	0.318
25	0.187	0.323	0.165	0.325
30	0.174	0.325	0.171	0.331
35	0.197	0.312	0.174	0.317
40	0.183	0.334	0.175	0.308
45	0.189	0.328	0.174	0.309
50	0.198	0.314	0.183	0.305
	Market condition category 2			
No. hedgers	PT	CoC	PT+HP	CoC+HP
5	0.224	0.302	0.195	0.280
10	0.210	0.313	0.195	0.283
15	0.227	0.299	0.196	0.279
20	0.233	0.299	0.192	0.275
25	0.229	0.291	0.197	0.284
30	0.211	0.302	0.198	0.290
35	0.229	0.288	0.206	0.277
40	0.216	0.301	0.202	0.281
45	0.220	0.294	0.202	0.284
50	0.214	0.297	0.202	0.287
	Market condition category 3			
No. hedgers	PT	CoC	PT+HP	CoC+HP
5	0.170	0.347	0.156	0.327
10	0.160	0.342	0.153	0.345
15	0.174	0.332	0.160	0.333
20	0.173	0.352	0.156	0.319
25	0.184	0.310	0.169	0.338
30	0.172	0.331	0.174	0.324
35	0.199	0.303	0.184	0.313
40	0.181	0.331	0.179	0.309
45	0.184	0.311	0.191	0.315
50	0.210	0.305	0.185	0.300

parent pool and probably make evolution less effective. It would therefore make sense to try increasing the number of evolutionary phases as the number of speculators in the simulation decreases; however, this could introduce other issues and would make it even harder to be sure the analysis of results is comparing like-with-like.

Beyond the scope of this paper, there are many things that could be done to make the simulation more sophisticated. Multiple markets could be simulated at once, perhaps alongside risk-free assets. Agents could then have a much more involved decision-making process, trading a portfolio of assets that varies as market conditions change. Transaction costs would also need to be much more realistic.

In summary, the intent here was to create an agent-based evolutionary simulation to allow an investigation of hedging pressure. The futures market produced by the simulation is significantly more realistic than previous work, and provides a good base for further work in that area. Previous analysis of futures has found evidence of hedging pressure in historical price data, and the experiments performed with the simulation indicate that hedging pressure does indeed exist. However, the evolutionary framework has demonstrated that when many competing agents make trades based on price predictions that take account of hedging pressure, any possible benefit disappears. The ability to see this feedback between agents goes beyond any analysis that could be performed using historical data. Hedging pressure does exist, but for real-world traders, trying to take account of it is unlikely to be worth the effort.

REFERENCES

[1] R. Alexrod and L. Tesfatsion. *A guide for newcomers to agent-based modelling in the social sciences*, volume 2. Elsevier, 2006.

[2] W. B. Arthur, J. H. Holland, B. LeBaron, R. G. Palmer, and P. Tayler. Asset pricing under endogenous expectations in an artificial stock market. In W. B. A. D. A. Lane and S. N. Durlauf, editors, *The Economy as an Evolving Complex System II*, Santa Fe Studies in the Sciences of Complexity Lecture Notes. Addison-Wesley, 1997.

[3] H. Bessembinder. Systematic risk, hedging pressure, and risk premiums in futures markets. *Review of Financial Studies*, 5:637–667, 1992.

[4] D. Challet, A. Chessa, M. Marsili, and Y.-C. Zhang. From minority games to real markets. *Journal of Quantitative Finance*, 1:168–176, 2001.

[5] D. Challet and Y.-C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407–418, 1997.

[6] C. Chiarella and G. Iori. A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2:346–353, 2002.

[7] V. Darley, A. Outkin, T. Plate and F. Gao. Sixteenths or pennies? Observations from a simulation of the Nasdaqstock market. In *Computational Intelligence for Financial Engineering*, 151–154, 2000.

[8] M. de la Maza and D. Yuret. A futures market simulation with non-rational participants. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 325–330. MIT Press, 1994.

[9] F. A. de Roon, T. E. Nijman and C. Veld. Hedging Pressure Effects in Futures Markets *Journal of Finance*, LV(3):1437–1456, 2000.

[10] J. D. Farmer. Toward agent-based models for investment. In *Developments in Quantitative Investment Models*, AIMR Conference Proceedings, pages 61–71, 2001.

[11] J. C. Hull. *Options, Futures and Other Derivatives*. Pearson Prentice Hall, 6th edition edition, 2006.

[12] R. Jagannathan. An investigation of commodity futures prices using the consumption-based intertemporal capital asset pricing model. *Journal of Finance*, 40:175–191, 1985.

[13] A. J. King, O. Streltchenko, and Y. Yesha. Using multiagent simulation to understand trading dynamics of a derivatives market. *Annals of Mathematics and Artificial Intelligence*, 44(3):233–253, 2005.

[14] B. LeBaron. Building the Santa Fe Artificial Stock Market. Available at http://people.brandeis.edu/ blebaron/wps.html, 2002.

[15] B. LeBaron. Agent-based computational finance. In L. Tesfatsion and K. Judd, editors, *Handbook of Computational Economics*, volume 2 of *Handbooks in Economics*, chapter 9. Elsevier, 2006.

[16] M. Roberto, S. Cincotti, S. M. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299(1-2):319–327, 2001.

[17] O. Streltchenko, Y. Yesha, and T. Finin. Multi-agent simulation of financial markets. In S. O. Kimbrough and D. J. Wu, editors, *Formal Modeling in Electronic Commerce: Representation, Inference, and Strategic Interaction*. Springer-Verlag, 2004.

[18] C. Wang. The behavior and performance of major types of futures traders. *Journal of Futures Markets*, 23(1):1–31, 2003.