ON THE NATURE OF THE STOCK MARKET: SIMULATIONS AND EXPERIMENTS

by

Hendrik J. Blok

B.Sc., University of British Columbia, 1993
M.Sc., University of British Columbia, 1995

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE STUDIES

(Department of Physics and Astronomy)

We accept this dissertation as conforming
to the required standard

____

____

____

____

____

THE UNIVERSITY OF BRITISH COLUMBIA

November 2000

Chapter 3

Decentralized Stock Exchange Model

3.1 Inspiration

I was growing disillusioned with CSEM when a collaboration with—then undergraduate student—Casey Clements inspired me to consider a radically different approach. Casey also expressed dissatisfaction with the price being set by a centralized control (market maker) and wondered how real markets worked. I explained to him that the stock (ticker) price was simply the last price at which a trade had occurred. Casey expressed an interest in modeling this approach but I explained to him the hurdles: namely that the agents would have to be a great deal more complicated than in current models because they would have to make complicated decisions involving two or more parameters.

Previously, agents were simple utility maximizers, applying Eq. 2.11 in order to calculate the quantity of shares they wanted to trade in response to a given price (set by the centralized control) but these new agents would have to decide on both the volume to trade and the price they wanted to trade at.

Nevertheless, Casey was enthusiastic so I obliged him and we constructed a simple event-driven model of stock exchange (neglecting the difficulties with the agents' decision processes) with the following properties:

1. agents trade by calling out and replying to orders,
2. trades can be called at any time (continuous time), and
3. agents can choose both the volume of the trade and the price.

In this model prices would be decentralized, arising directly from the agents decisions rather than being governed by a market maker. Thus was born the *Decentralized Stock Exchange Model* (DSEM).

3.2 Basic theory

This model is of a more original nature than CSEM was and, therefore, requires more explanation. For this reason the theory is divided into two sections, one which explains the basic structure of the model and another which describes how fluctuations are incorporated. First, the structure of the model will be developed.

3.2.1 Assumptions

The decentralized model discussed here contains many of the same assumptions as CSEM (heterogenous agents; the market is composed of a single risky asset and a single riskless asset; cash and shares are both conserved; et cetera). For the sake of brevity, only the differences in assumptions between the two models will be discussed.

Decentralization

The primary difference between the two models is, obviously, the move to a decentralized market. This means that the agents are allowed to trade directly with one another without any interference from a market maker or specialist. The market maker may be interpreted as having been relegated to the mechanical role of matching buyers with sellers, with no influence on the price. This interpretation resembles intra-day trades of a fairly active stock [26]. CSEM, in contrast, was meant to mimic low-frequency trading, on timescales no shorter than a single day.

Continuous time with discrete events

By moving to intra-day trading, the natural periodicity of the market opening and closing daily, which gave rise to discretized time in CSEM, is eliminated. In fact, in DSEM it is assumed that the market never closes; it is open around the clock, 24 hours per day. Trades may be executed at any instant and time is a continuous variable. (Another interpretation is that the market does close but when it re-opens it continues from where it left off without any effect from the close.)

To implement continuous-time (at least to some fine resolution) on fundamentally discrete devices (digital computers) a shift of paradigms is required. Traditionally, time evolution is simulated by simply incrementing "time" by a fixed amount (as in CSEM). This approach is cumbersome and inefficient when the model consists of discrete events (eg. trades) occurring at non-uniform time intervals. Instead, an event-driven approach is preferred [54] in which waiting times (delays) for all possible events are calculated and "time" is immediately advanced to the earliest one. (As each event time is calculated the event is placed in an ordered queue so the earliest event is simply the first event in the queue.) The basic algorithm follows:

1. Calculate waiting times for all events.
2. Find event with shortest waiting time.
3. Advance time to this event and process it.
4. Recalculate waiting times as necessary.
5. Return to Step 2.

Separation of time scales

The model allows two types of events: *call* orders and *reply* orders. By distinguishing between the two a simplifying assumption can be made: no more than one call order is active at any time.

In many real markets orders are good (active) until filled or until they expire. When new orders are placed they are first treated as reply orders by checking if they can satisfy any outstanding orders, and then, if they haven't been filled, they are placed on a *auction book*, and become call orders until they are removed.

In DSEM, however, it is assumed that reply orders occur on a *much* faster timescale than calls. As soon as a call order is placed, all reply orders are submitted and executed (almost) instantaneously and the call (if not filled) can immediately be expired (because no more replies are expected). Hence, the probability that two (call) orders are active simultaneously becomes negligible and it is assumed that at most one is ever active at any given time.

This assumption improves performance at the cost of realism. Allowing multiple orders to be simultaneously active would require more complicated bookkeeping and would degrade simulation performance. (It was observed that DSEM exhibited rich enough behaviour that the assumption did not need to be discarded.)

3.2.2 Utility theory

As with CSEM, we again begin with a game theoretic approach. However, instead of an exponential utility function, a power-law is used,

$$U(w) = \left\{ \begin{array}{l|l} \frac{1}{1-a} w^{1-a} & a \neq 1 \\ \ln w & a = 1. \end{array} \right. \tag{3.1}$$

Notice these forms are effectively identical as $a \to 1$ because $\lim_{a \to 1} U' = w^{-1}$ for both, and utility is only defined up to a (positive) linear transformation.

The advantage of the power-law utility (sometimes known as the Kelly utility [55]) is that it has a constant *relative* risk aversion

$$R(w) \equiv -\frac{wU''}{U'} = a, \tag{3.2}$$

unlike the exponential utility which, from Eq. 2.4, has constant *absolute* risk aversion $A(w) \equiv -U''/U' = 1/w_{goal}$. The constant relative risk aversion eliminates the scaling problems which had to be worked around in Section 2.2.4 and keeps an agent equally cautious regardless of its absolute wealth. (See Ref. [56] for a more detailed description of absolute and relative risk aversion).

3.2.3 Optimal investment fraction

Let us assume an agent is interested in holding a fixed fraction i of its capital in the risky asset. If the price moves from $p(0)$ to $p(t)$ the return-on-investment is

$$r(t) = \frac{p(t) - p(0)}{p(0)} \tag{3.3}$$

and the final wealth can be written as

$$\begin{aligned} w(t) &= w(0)(1-i) + w(0)i(r+1) && (3.4) \\ &= w(0)(1+ir). && (3.5) \end{aligned}$$

The goal is to find the value of i which maximizes the expected utility at some future time t. First, the expected utility must be calculated:

$$\langle U \rangle = \frac{w(0)^{1-a}}{1-a} \left\langle (1+ir)^{1-a} \right\rangle. \tag{3.6}$$

Unfortunately, finding a closed-form solution for $\langle U \rangle$ is difficult, even with very simple probability distributions.

However, if we assume the timescale is relatively short then we expect the returns to be small and the above equation can be expanded around $r = 0$, as

$$\langle U \rangle \approx w(0)^{1-a} \left\langle \frac{1}{1-a} + ir - \frac{1}{2} a i^2 r^2 + \cdots \right\rangle . \tag{3.7}$$

Thus, the optimization condition becomes

$$\begin{aligned} 0 &= \frac{d \langle U \rangle}{di} && (3.8) \\ &= w(0)^{1-a} \left[\langle r \rangle - a i^* \left\langle r^2 \right\rangle \right] && (3.9) \\ &= w(0)^{1-a} \left[\langle r \rangle - a i^* (\mathrm{Var}\,[r] + \langle r \rangle^2) \right], && (3.10) \end{aligned}$$

giving an optimal investment fraction

$$i^* = \frac{\langle r \rangle}{a(\mathrm{Var}\,[r] + \langle r \rangle^2)} \tag{3.11}$$

with the constraints $0 \leq i^* \leq 1$. (This derivation implicitly assumed that the higher moments in the expansion are negligible—an assumption which may not be valid even for short timescales, as will be seen in Chapter 5).

Note how closely this corresponds with Eq. 2.19 in CSEM, differing only by an extra term in the denominator. However, the use of a power-law utility allowed us to drop many of the assumptions originally required, such as explicitly hypothesizing that the returns were Gaussian-distributed.

3.2.4 Fixed investment strategy

Eq. 3.11 states that, given some expectation and variance of the future return, one should hold a constant fraction of one's capital in the risky asset. The same result was derived by Merton [56, Ch. 4] assuming a constant consumption rate (which can be taken to be zero, as in DSEM). Further, Maslov and Zhang [57] demonstrated that keeping a fixed fraction of one's wealth in the risky asset maximizes the "typical" long-term growth rate (defined as the *median* growth rate).

These references suggest that the optimal strategy is to keep a fixed fraction of one's wealth in stock—a Fixed Investment Strategy (FIS). The FIS is empirically tested with a hypothetical portfolio in Chapter 6 and developed, here, for use with DSEM. The strategy discards the attempt to forecast returns (which proved problematic in CSEM) in favour of the basic principle of maintaining some fixed fraction invested in stock.

Ideal price

The total value of a portfolio consisting of c cash and s shares at price p is

$$w(p) = c + ps. \tag{3.12}$$

The goal of FIS is to maintain a balance

$$ps = iw(p) \tag{3.13}$$

by adjusting one's holdings s. (The ideal investment fraction is simply denoted by i henceforth, without the cumbersome asterisk.) If an agent currently holds c cash and s shares then the agent would achieve the optimal investment fraction i if the stock price was

$$\boxed{p^* = \frac{ic}{(1-i)s},} \tag{3.14}$$

which will be called the agent's *ideal price.*

If the price is higher than the ideal price then the agent holds too much capital in the form of stock and will want to sell and, conversely, at a lower price the agent will buy in order to buoy up its portfolio.

Optimal holdings

The fixed investment strategy specifies not only what type of order to place at a given price, but also precisely how many shares to trade. At a price p an agent would ideally prefer to hold stock

$$\boxed{s^* = \frac{iw(p)}{p}.} \tag{3.15}$$

Given the ideal price above, the agent's wealth can be written

$$w(p) = \frac{1-i}{i}p^*s + ps. \tag{3.16}$$

Hence, an agent's optimal trade at price p is

$$\begin{aligned} s^* - s &= \frac{iw}{p} - s && (3.17)\\ &= (1-i)s\frac{p^*}{p} + is - s && (3.18)\\ &= (1-i)\left(\frac{p^*}{p} - 1\right)s. && (3.19) \end{aligned}$$

The fractional change of stock $(s^* - s)/s$ is plotted as a function of price in Fig. 3.1.

So, given a portfolio (c,s) and an optimal investment fraction i, FIS prescribes when to buy and sell and precisely how much.

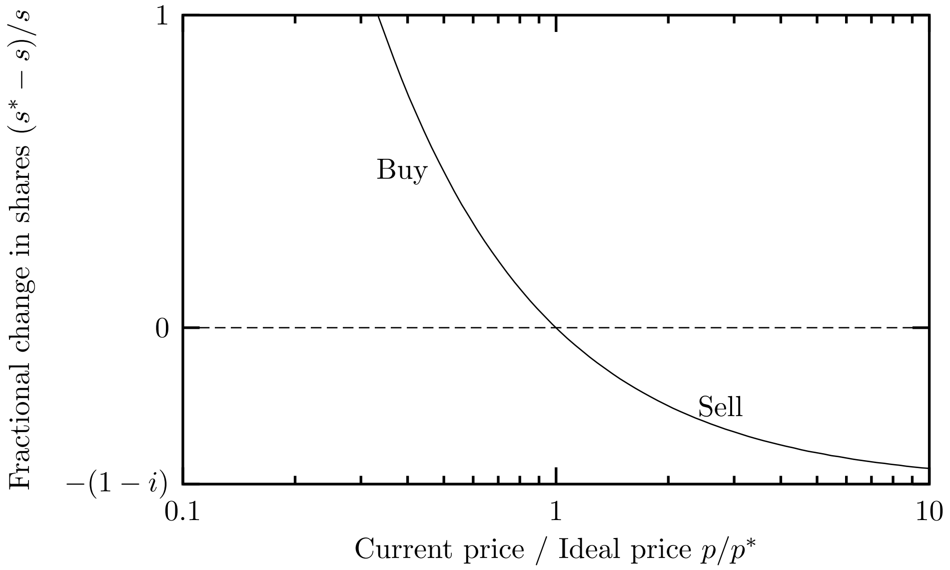

Figure 3.1: The fixed investment strategy specifies how many shares to trade at a price p given an investment fraction i and an ideal price p^* (from Eq. 3.14). As the current price drops toward zero the fractional change in shares diverges.

3.2.5 Friction

The derivations of the fixed investment strategy here and elsewhere [56, 57] assume no transaction costs, which makes it impractical in real markets. Every minuscule fluctuation of a stock's price would require a trade in order to rebalance the portfolio but if the fluctuation was too small the trade would cost more than the value of the shares traded, resulting in a net loss.

To circumvent this problem in a simulated portfolio with a commission on every trade (see Chapter 6) I imposed limits on the buying and selling prices:

$$p_B = p^*/(1+f) \tag{3.20}$$

$$p_S = p^*(1+f) \tag{3.21}$$

where f is defined as the trading *friction*.

The same approach can be used here: don't buy until the price drops below the limit p_B and don't sell until it rises above the limit p_S. This allows the simulation to mimic the effect of transaction costs while conserving the total cash. It is also required to make the simulation a discrete-event model. Without it, agents would trade on a continuous-time basis and the model would be difficult to simulate (and less realistic).

Obviously, the larger the friction, the larger a price fluctuation will be required before an agent decides to trade. So increasing friction decreases the trade frequency and hinders market activity—which explains why the term "friction" is used.

Each agent's friction must be strictly positive because a zero value would allow an agent to place an order to trade zero shares at its ideal price—a null order. The simulation does not forbid this but, as will be seen, such an order would never be accepted.

There is no fixed upper limit on the friction but it seems reasonable to impose $f < 1$. This would mean an agent places buy and sell limits at one half and double the ideal price, respectively. (It would be peculiar for an investor not to sell when a stock's price doubles!)

Thus, the friction f is the first agent-specific parameter in DSEM and it is chosen such that $f \in (0, 1)$.

3.2.6 Call orders

As mentioned above, DSEM is a discrete-event simulation. The events are orders which are called out. An order consists of a price p_o, type ("Buy" or "Sell") and

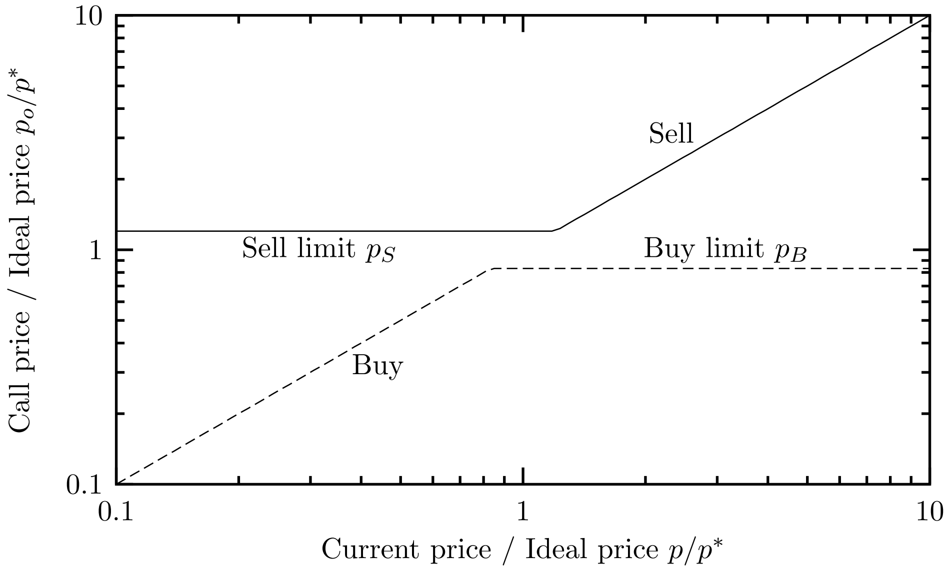

Figure 3.2: Call orders p_o are placed at either the current price p or the limit prices p_B or p_S, whichever is better. The spread between the limits increases with friction f.

quantity of shares to trade.

The agents use the limit prices discussed in the last section to set their trading prices, thereby solving the problem of how to design agents which can choose both a trade volume *and* price. (The volume is set by Eq. 3.19.) Of course, if the current market (last traded) price p is "better" than the limit price ($p > p_S$ or $p < p_B$) then the agents rationally choose to trade at that price:

$$\boxed{p_o = \begin{cases} \min(p, p_B) & \text{for "Buy" orders} \\ \max(p, p_S) & \text{for "Sell" orders.} \end{cases} \qquad (3.22)}$$

These order prices, shown in Fig. 3.2, are substituted into Eq. 3.19 to compute the volume of shares to trade. Notice that there are two options (buy or sell) for every price. Before discussing how the agent decides which action to take, we must understand the process which governs most discrete-event simulations.

3.2.7 Poisson processes

A Poisson process is a stochastic counting process defined by the probability of an event occurring in an infinitesimal interval dt,

$$\pi(dt) = dt/\tau \tag{3.23}$$

where τ is the average event interval [58]. All events are assumed to be independent. The cumulative probability of *no* events having fired within a finite interval t is then given by

$$\Pr(0, t) = e^{-t/\tau} \tag{3.24}$$

which accurately describes many natural processes, such as radioactive decay of nuclei.

The advantage of the Poisson process from a simulation perspective is that it may be interrupted. Consider a time interval t which is divided into two intervals t_1 and t_2 ($t = t_1 + t_2$). The probability of no events within t can be written

$$\Pr(0, t) = e^{-(t_1+t_2)/\tau} \tag{3.25}$$

$$= \Pr(0, t_1)\Pr(0, t_2) \tag{3.26}$$

which simply means that if the event did not occur within the interval t then it must not have occurred within either of the sub-intervals t_1 or t_2.

This equivalence relation may be interpreted to mean that a clock which measures the stochastic time to a Poisson event may be reset at any time before the event fires, without changing the probability distribution. This property is very useful for discrete-event simulations because it allows one to proceed to the first event time, update the system, and recalculate all the event times from this point without disturbing the process. (Event times would need to be recalculated, for instance, if their average rate τ was affected by the event which transpired.)

We now have the background necessary to discuss how an agent chooses which action to take.

3.2.8 Call interval

Agents always have two options when deciding on a call order: place a "Buy" or a "Sell" order. Since the simulation is event-driven, the easiest way to handle this is to allow both possibilities. Each is an event which will occur at some instant in time. The events of calling orders is modeled as a Poisson process, as discussed above.

More precisely, "Buy" and "Sell" orders are modeled as independent Poisson processes, each with its own characteristic rate. Intuitively, if the current stock price

p is very high then we expect agents will try to sell at or near the current price, rather than trying to buy at a much lower price, and vice versa if the price is low. This intuition can be captured by making the Poisson rates for the "Buy" and "Sell" calls price dependent [26], as

$$\boxed{\begin{aligned} \tau_B(p) &= \frac{p}{p_B}\tau && (3.27) \\ \tau_S(p) &= \frac{p_S}{p}\tau && (3.28) \end{aligned}}$$

where τ is an unspecified (constant) timescale and τ_B and τ_S are the average times between each type of call order.

The above linear price dependence is justified only on the basis of its simplicity, but it also has some reasonable consequences. An agent makes the fewest calls (of either type) when the call rate $1/\tau_{B+S}$ is minimized,

$$0 = \frac{d}{dp}\frac{1}{\tau_{B+S}} \tag{3.29}$$

$$= \frac{d}{dp}\left(\frac{1}{\tau_B} + \frac{1}{\tau_S}\right) \tag{3.30}$$

$$= -\frac{p_B}{p^2} + \frac{1}{p_S} \tag{3.31}$$

which occurs at a price $p = \sqrt{p_B p_S} = p^*$, the agent's ideal price. Hence, when an agent is satisfied with its current portfolio, it will place the fewest (speculative) orders.

As the price moves away from the agent's ideal price the call rate increases reflecting an increased urgency. With many agents independently placing call orders, the agent with the greatest urgency (shortest waiting time) will tend to place the first order, so urgency drives market fluctuations [59]. This is in contrast with CSEM and many other simulations in which fluctuations are driven by supply vs. demand [27–36].

Incidentally, the minimum call rate is

$$\frac{1}{\tau_{B+S}(p^*)} = \frac{1}{\tau_B(p^*)} + \frac{1}{\tau_S(p^*)} \tag{3.32}$$

$$= \frac{2}{(1+f)\tau} \tag{3.33}$$

which decreases as the friction f is increased, as discussed previously.

Eqs. 3.27–3.28 were introduced to increase the probability of buying if the last trading price was low and selling if the last trading price was high. The probability

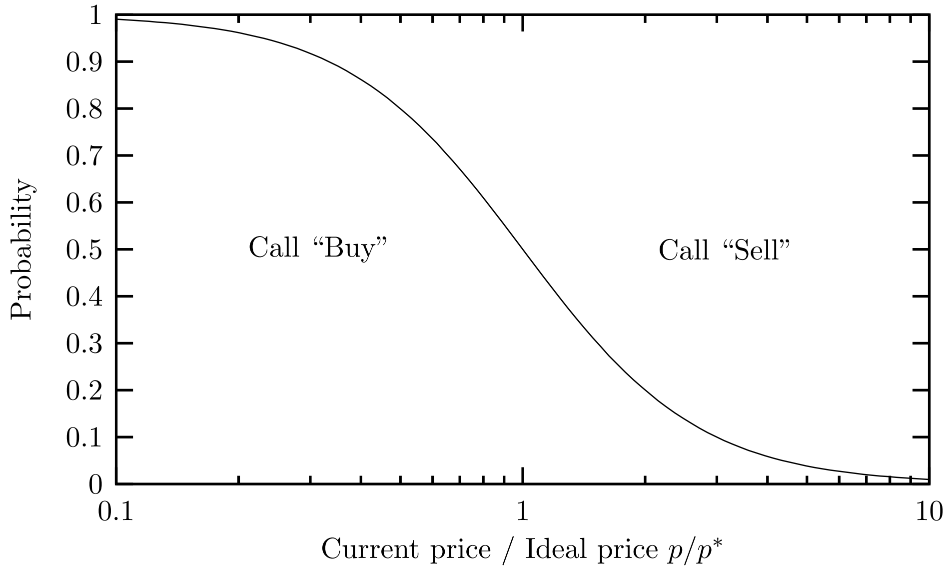

Figure 3.3: "Buy" and "Sell" call orders are modeled as independent Poisson processes with price-dependent rates. As the last trading price increases, the probability of a "Sell" order being called becomes much more likely than a "Buy."

of the next call being a "Buy" can be explicitly calculated: consider the probability of a Buy order being placed between times t and $t + dt$ with neither a Buy nor a Sell having occurred yet. Then, marginalizing over all t gives the probability of a Buy occurring first,

$$\begin{aligned}
\Pr(\text{Buy}) &= \int_0^\infty \frac{dt}{\tau_B} e^{-t/\tau_B - t/\tau_S} && (3.34)\\
&= \frac{1/\tau_B}{1/\tau_B + 1/\tau_S} && (3.35)\\
&= \frac{p^*/p}{p^*/p + p/p^*}. && (3.36)
\end{aligned}$$

Of course, the probability of a Sell order being placed first is the complement, $\Pr(\text{Sell}) = 1 - \Pr(\text{Buy})$. Fig. 3.3 shows that Buy orders are much more likely at low prices and Sell orders more likely at high.

3.2.9 Reply orders

Orders are divided into two types: call orders and reply orders. The call orders have been discussed above. Reply orders are handled in much the same manner with two important changes:

1. When replying to a call order, the replier is not free to set a price but must accept the called price.
2. Reply orders happen on a much faster timescale than call orders.

The first item requires that reply orders be handled slightly differently than call orders. The agent still calculates price limits according to Eqs. 3.20–3.21 but now this is used as the criterion for whether to place an order or not. If the order price meets the Sell limit $p_o \geq p_S$ then a "Sell" reply is placed and if the price meets the Buy limit $p_o \leq p_B$ then a "Buy" reply is placed. Otherwise the agent does not reply to the called order.

It is assumed that a called order is completely transparent; all potential repliers have full information about the order including the price, type of order (Buy or Sell) and quantity of shares to be traded. This provides another criterion whether to reply or not: the quantity of shares in the called order must be sufficient to completely fill the replier's demand. Although this requirement may seem too strict, it is useful because it prevents callers from swaying the price series with negligible orders (as could occur with zero friction or if the caller has negligible wealth compared to the replier).

Like call orders, replies receive stochastic execution times with average intervals from Eq. 3.27 or Eq. 3.28 depending on whether the reply is a Buy or Sell order. Replies compatible with the call order are then processed in a "first come, first served" queue until the call order is filled or all repliers are satisfied (partial orders are processed).

As mentioned above, replies occur on a much faster timescale than call orders, such that no more than one call order is ever active at a time. The time for all replies to be processed is assumed to be infinitesimal compared to the calling time interval.

3.2.10 Time scale

If each agent places Buy and Sell call orders at rates $1/\tau_B$ and $1/\tau_S$, respectively, then the net rate ρ of call orders being placed (for all agents) is

$$\rho = \sum_{i \in \text{agents}} \left(\frac{1}{\tau_{B,i}} + \frac{1}{\tau_{S,i}} \right). \tag{3.37}$$

Let us assume that all agents are satisfied with their current portfolios; an assumption which, by Eq. 3.33, minimizes the net event rate,

$$\rho_{min} = \frac{2}{\tau} \sum_i \frac{1}{1 + f_i}. \tag{3.38}$$

Now we can identify the minimum event rate with a real rate in order to specify the timescale τ. Note that fixing τ is only necessary in order to set a scale for comparing simulation data with empirical market data. An arbitrary but convenient choice is to assume each agent trades once each day (on average). Then, if the market contains N agents, we expect $\rho_{min} = N/2$ (because each call order is a trade between at least two agents) so the timescale τ should be

$$\boxed{\tau = \frac{4}{N} \sum_i \frac{1}{1 + f_i}.} \tag{3.39}$$

Hence, one time unit is meant to represent one day. This is a particularly useful choice because it allows us to draw some parallels between DSEM and the original model, CSEM, in which agents were constrained to exactly one trade per day. Notice the parallel is not perfect for two reasons: firstly, there is no guarantee that a call order will have any repliers and secondly, if it is executed there may be multiple repliers. Nevertheless, the scaling should be accurate within an order of magnitude.

3.2.11 Initialization

DSEM is initialized with N agents amongst whom some total cash C and total shares S are distributed. (N sets the "size" of the market. Heavily traded stocks would be represented by large N.) As in CSEM, the cash and shares will usually be distributed uniformly between the agents.

Each agent begins with an ideal investment fraction $i(0) = 1/2$ which gives an ideal price

$$p^*(0) = \frac{c}{s}. \tag{3.40}$$

For simplicity, each agent begins satisfied with its current portfolio, believing that the current market price is actually its ideal price $p(0) = p^*(0)$.

Notice that Eq. 3.11 is not used to set the investment fraction. Its purpose was only to establish that holding a fixed fraction i of one's wealth in stock is rational. How i is updated is the subject of the next section.

If the cash and shares are initially distributed equally amongst the agents (as will be assumed for all runs, unless otherwise stated) then the simulation begins in stasis: any Sell order submitted must necessarily be set above all repliers ideal price (C/S) so it will not be filled, and vice versa for Buy orders. Thus, no trading will occur and the price will never move away from C/S. What is needed is some stochastic driving force to initiate the dynamics. (Even starting with a non-uniform distribution produces only transient fluctuations before the price stabilizes.)

In this section three market parameters were introduced: the number of agents N, the total cash C, and the total shares S.

3.3 Fluctuation theory

The above theory completely specifies the basic model. What remains is to incorporate fluctuations, as discussed in this section. Recall that forecasting was problematic in CSEM (see Fig. 2.9, for example) so in DSEM a different tack is taken.

3.3.1 Bayes' theorem

Before we can understand how fluctuations are incorporated it is necessary that we briefly review some results from Bayesian probability theory [25, Ch. 4], which provides an inductive method for updating one's estimated probabilities of given hypotheses as new data arrive:

Let H represent some hypothesis which one wishes to ascertain the truth value of. If X is our prior information then we begin with a probability of H given

X, denoted by $P(H|X)$. Now notice that as new information D (data) arrives the joint probability of both the hypothesis *and* the data becomes

$$P(HD|X) = P(H|X)P(D|HX) \tag{3.41}$$
$$= P(D|X)P(H|DX) \tag{3.42}$$

from the product rule of probabilities.

But these equations can be rewritten to give the probability of the hypothesis in light of the new information,

$$P(H|DX) = P(H|X)\frac{P(D|HX)}{P(D|X)}, \tag{3.43}$$

which is known as Bayes' theorem.

Let us define *evidence* as the logarithm of the odds ratio $e \equiv \log\left[P/(1-P)\right]$, which is just a mapping from probability space to the set of all real numbers $(0,1) \rightarrow (-\infty, \infty)$.

The advantage of the evidence notation over probabilities is that incorporating new information is an additive procedure

$$e(H|DX) = e(H|X) + \log\left[\frac{P(D|HX)}{P(D|\bar{H}X)}\right], \tag{3.44}$$

where $\bar{H}$ is the negation of H.

With the understanding that assimilating new information is an additive process for evidence, we may proceed with extending DSEM to include fluctuations.

3.3.2 News

In real markets a stock's price is derived from expectations of its future earnings. These expectations are formed from information about the company, which is released as *news.* In other words, news drives market fluctuations.

In DSEM fluctuations are also driven by news. How to represent news in this model, though, is problematic. Inspiration comes from Cover and Thomas [24, Ch. 6] in which the optimal (defined as maximizing the expect growth rate of wealth) wagering strategy in a horse race (given fair odds) is to wager a fraction of one's capital on each horse equal to the probability of that horse winning—a fixed investment strategy, where the investment fraction is identified with a probability.

In DSEM, this would have to be translated as the probability of the stock (or cash) "winning," the interpretation of which is unclear. (A loose interpretation might be that the stock wins if its value at some future horizon is greater than cash, otherwise cash wins.)

Regardless of what i is a measure of, if we interpret i as *some* probability measure then Bayesian probability theory offers an avenue for further development. The evidence, as discussed in the last section, corresponding to the probability i is

$$\boxed{e = \log \frac{i}{1-i}. \qquad (3.45)}$$

As new information is acquired the evidence is updated via

$$e' = e + \eta \qquad (3.46)$$

where η represents the complicated second term of Eq. 3.44.

This suggests modeling news as a stochastic process η which affects an agent's confidence in the stock and, thereby, investment fraction. (Modeling the news as η directly, instead of through the information D as presented in Eq. 3.44, dramatically simplifies the calculations.) A positive η increases the evidence (and investment fraction), a negative value decreases it, while $\eta = 0$ is neutral.

Assuming news releases have a finite variance and are cumulative, the Central Limit Theorem indicates the appropriate choice is to model η as Gaussian noise. It should have a mean of zero $\langle \eta \rangle = 0$ (unbiased) so that there is no long-term expected trend in investment (or price).

The scale of the fluctuations σ_η, though, is difficult to decide; but this just opens up an opportunity to increase diversity amongst the agents—let each set its own scale. We begin by setting an arbitrary scale σ_η and then allowing agents to rescale it according to their own preferences. Since each agent will apply its own scaling factor the universal scale σ_η is arbitrary so it will be set to a convenient value later.

News response

First, we begin by defining individual scale factors. Let us define a new agent-specific parameter r_n, which represents the agent's *responsiveness* to news. Then evidence would be updated as

$$\boxed{e' = e + r_n \eta. \qquad (3.47)}$$

A responsiveness of zero would indicate the agent ignores news releases and maintains a constant investment ratio. As the responsiveness increases the agent becomes increasingly sensitive to news and adjusts its ideal investment fraction more wildly, via

$$i' = \frac{i \exp(r_n \eta)}{1 - i\,(1 - \exp(r_n \eta))}. \qquad (3.48)$$

Notice there exists a symmetry between positive and negative news if the responsiveness also changes sign, $(-r_n)(-\eta) \equiv r_n\eta$, so we can impose the restriction $\langle r_n \rangle \geq 0$ (averaged over all agents) without loss of generality.

News releases

Since DSEM is a discrete-event simulation the news must also be inserted as discrete events. For reasons discussed in Appendix B news events are modeled as a discrete Brownian process with some characteristic interval τ_n (unbiased Gaussian-distributed jumps occurring regularly at intervals of τ_n). Thus, on average, we expect $1/\tau_n$ news events each day. If every news release has a variance of σ_η^2 then the variance of the cumulative news after one day is σ_η^2/τ_n. To minimize the impact of the news interval parameter the variance of the news over some fixed interval should be constant, independent of how often news is released. Otherwise rescaling the news interval will rescale the evidences and hence impact upon the scale of the price series. Let us take the variance to be one unit over one day, which is satisfied when

$$\sigma_\eta^2 = \tau_n. \tag{3.49}$$

To draw parallels with real markets the news release interval is chosen between $1/6.5 \leq \tau_n \leq 5$ where the market is assumed to be open six and one half hours per day [7]. Thus news releases occur at least once per week and at most once per hour. A smaller interval is inappropriate because news is irrelevant until an investor is made aware of it—most investors (except professional traders) probably do not check for news more than once per hour. (In fact, most people still get their news from the daily newspaper, suggesting $\tau_n = 1$.) On the other hand, a timescale longer than a week (5 days) is useless because we are particularly interested in fluctuations on the scale of hours to days, so the driving force should be on the same scale.

Parameters

In this section two new parameters related to news were presented: the agent-specific news response parameter r_n which is allowed to be negative but is constrained such that $\langle r_n \rangle \geq 0$; and the global news interval parameter which is strictly positive and constrained to $\tau_n \in (1/6.5, 5)$.

3.3.3 Price response

DSEM with news-driven fluctuations is a complete model ready for experimentation. However, it is lacking in that it neglects a significant source which real investors often construct their expectations from: the price history itself. It is probable that feedback from the price series is integral to the clustered volatility and other complex phenomena found in empirical data. In CSEM this feedback was modeled by tracking the history of returns and forming future expectations therefrom, to set the investment fraction. Since DSEM sets the investment fraction completely differently, this method is unavailable. However, it is possible to construct a method which allows agents to extract information from the price series.

Consider how a single agent's ideal price is affected by news. From Eq. 3.14, after a news release η the ideal price becomes

$$\begin{aligned} p^{*\prime} &= \frac{i' c}{(1-i')s} && (3.50) \\ &= \exp(r_n \eta) p^* && (3.51) \end{aligned}$$

which suggests that news is related to the logarithm of the price through

$$\eta \propto \log \frac{p'}{p}. \tag{3.52}$$

Therefore price movements imply news and if the price does change, an agent may infer that it missed some news which others are privy to. Thus the price feedback may be inserted by extending the evidence dependence such that on a price move from p to p'

$$\boxed{e' = e + r_p \log \frac{p'}{p}} \tag{3.53}$$

where r_p, the *response to price*, is a new agent-specific parameter. Setting $r_p = 0$ eliminates the price feedback and reverts the model to being driven solely by news. But with non-zero r_p agents have a *chartist* nature: they presume the price series contains information (trends) and wager accordingly.

Some market models separate agents into two groups: *fundamentalists* and *chartists* [32, 36, 47]. Fundamentalists (or "rational" traders) value the stock using fundamental properties such as dividends and reports of assets. In the absence of dividends, this would correspond to responding strongly to news releases in DSEM. Chartists (or "noise" traders) simply use the price history itself as an indicator of the stock's value. This would correspond to a strong price response in DSEM. However, in DSEM the two are not exclusive. Instead of drawing a distinction between the two types of traders a continuum exists where r_n and r_p can take on a wide range of values so that agents may value the stock based on fundamentals *and* on its performance history.

Note that the price only changes when a trade occurs. Only agents not involved in the trade (neither the caller nor a satisfied replier) should update their evidence since the traders can't be interpreted as having "missed" some information (because their trade *was* the information). Further, allowing the caller and replier(s) to also update there evidence would cause complications because an agent could never reach equilibrium—a trade would bring them to their ideal investment fractions but then their evidences (and investment fractions) would be immediately changed.

The price response parameter is similar to an autocorrelation in returns. Therefore we expect the dynamics to destabilize if we allow it to exceed unit magnitude. As we will see in Chapter 4, imposing $-1 < \langle r_p \rangle < 1$ keeps the price from rapidly diverging.

In this section a new parameter, the price response r_p which is constrained by $|\langle r_p \rangle| < 1$, was derived.

3.3.4 Review

The Decentralized Stock Exchange Model (DSEM) consists of a number N of agents which trade with each other directly, without the intervention of a market maker. In this section the structure of the model will be reviewed.

Game theory indicates the optimal strategy is to maintain a fixed fraction of one's capital in stock, the Fixed Investment Strategy (FIS). Therefore the agents trade in order to rebalance their portfolios consisting of c cash and s shares. News releases and price fluctuations cause agents to re-evaluate their investment fraction i and up- or down-grade it as they see fit.

Algorithm

Everything that happens in the model occurs as a discrete event in continuous time. The basic algorithm follows:

1. Initialization. Cash and shares distributed amongst agents. Agents sample Poisson distribution to get waiting times until first call orders. Waiting time for first news event is set to zero (first event).

2. Next event is found (shortest waiting time). Time is advanced to that of the event which is executed next. If it is a news release proceed to Step 3. Otherwise it is a call order being placed, proceed to Step 4.

3. News event. Sample Gaussian distribution to generate deviate for news. Adjust all agents' ideal investment fractions. Set news waiting time to τ_n and recalculate all agents' waiting times. Return to Step 2.

Symbol	Interpretation	Range
	Market parameters	
N	number of agents	2+
C	total cash available	
S	total shares available	
τ_n	average interval between news releases	$(1/6.5, 5)$
	Market state variables	
$p(t)$	stock price at time t	
$v(t)$	trade volume (number of shares traded) at time t	
	Agent parameters	
f_j	friction of agent j	$(0, 1)$
$r_{n,j}$	news response of agent j	$\langle r_n \rangle \geq 0$
$r_{p,j}$	price response of agent j	$\lvert\langle r_p \rangle\rvert < 1$
	Agent state variables	
c_j	cash held by agent j	
s_j	shares held by agent j	
$w_j(p)$	wealth of agent j at stock price p	
i_j	optimum investment fraction of agent j	

Table 3.1: All parameters and variables used in the Decentralized Stock Exchange Model (DSEM).

4. Call order placed.
   (a) Calculate all reply orders and corresponding waiting times. Place compatible replies in queue, ordered by waiting times.
   (b) If queue is empty, proceed to Step 4d. If queue is not empty, remove first reply order from queue and execute it.
   (c) Reduce outstanding call order by appropriate volume. If not completely filled, return to Step 4b.
   (d) Recalculate call-order waiting times for agents which traded. If price did not change, return to Step 2.
   (e) If price did change, recalculate ideal fractions for all other agents. Recalculate call-order waiting times. Return to Step 2.

Parameters

For convenience all the variables used in DSEM are listed in Table 3.1. The parameters are inputs for the simulation and the state variables characterize the state of the simulation at any time completely. For each run, the agent-specific parameters are set randomly; they are uniformly distributed within some range (a subset of the ranges shown in the table). Each dataset analyzed herein will be characterized by listing the market parameters and the ranges of agent parameters used.

3.4 Implementation

The Decentralized Stock Exchange Model (DSEM) is completely characterized by the above theory. The model is beyond the scope of rigorous analysis in all but the most trivial of scenarios so it is simulated via computer. The model was programmed in C++ using Borland C++Builder 1.0 on an Intel Pentium II computer running Microsoft Windows 98. The source code and a pre-compiled executable are available for download from `http://rikblok.cjb.net/phd/dsem/`.

DSEM encounters some of the same issues as CSEM. In particular, random numbers are handled using the same code as was discussed in Section 2.3.1. Similarly, the random number seed will be specified with the other model parameters if the default (the current time) is not used.

3.5 Parameter space exploration

Having converted DSEM to computer code time series can be generated for numerical analysis. Currently DSEM requires seven parameters (one fewer than CSEM started with) to fully describe it. Again, it would be useful to try and reduce the parameter space before collecting any serious data.

3.5.1 Number of agents N

The effect of changing the number of traders will be explored in detail in Chapter 4 and is left until then.

3.5.2 Total cash C and total shares S

In this section the effect of rescaling the total cash C and total shares S will be explored. Let us denote rescaled properties with a prime. Then rescaling cash by a factor A and shares by B is written

$$C' = AC \tag{3.54}$$

$$S' = BS. \tag{3.55}$$

Parameter	Run 1	Run 2	Run 3
N	100	100	100
C	\$1,000,000	\$10,000,000	\$1,000,000
S	1,000,000	10,000,000	10,000,000
τ_n	1	1	1
r_n	0.025 ± 0.025	0.025 ± 0.025	0.025 ± 0.025
r_p	0.75 ± 0.75	0.75 ± 0.75	0.75 ± 0.75
f	0.05 ± 0.05	0.05 ± 0.05	0.05 ± 0.05
seed	-2	-2	-2

Table 3.2: Parameter values for DSEM Runs 1, 2 and 3.

Cash and shares are rescaled equally for each agent so the distribution remains constant.

We begin by noticing that the initial evidence is unchanged (being fixed at $e(0) = 0$). The evidence depends on news releases (which are unaffected by rescaling) and price movements through Eq. 3.53. Assuming price scales linearly ($p' \propto p$) the logarithm of the price ratios will be unchanged. Thus, the evidence and the investment fraction will also be unchanged under rescaling.

Assuming i remains unscaled, an agent's ideal price scales as

$$p^{*\prime} = \frac{ic'}{(1-i)s'} = \frac{A}{B}p^* \tag{3.56}$$

and the buy and sell limits scale identically. Since all agents ideal prices scale as A/B it is reasonable to expect that the entire price series scales as

$$p' = \frac{A}{B}p, \tag{3.57}$$

as is required for the investment fraction to remain unchanged. Thus, these two hypotheses are compatible and, through the initial conditions, are realized.

This argument took the same form as in Section 2.4.2 for CSEM. An identical argument for trade volume also applies giving the result that volume scales with shares as

$$v' = Bv. \tag{3.58}$$

To test these hypotheses three runs were performed, with the parameter values shown in Table 3.2. Contrasting Runs 2 and 3 with Run 1 give $A = 10, B = 10$ and $A = 1, B = 10$, respectively. The results, shown in Fig. 3.4, confirm our hypothesis.

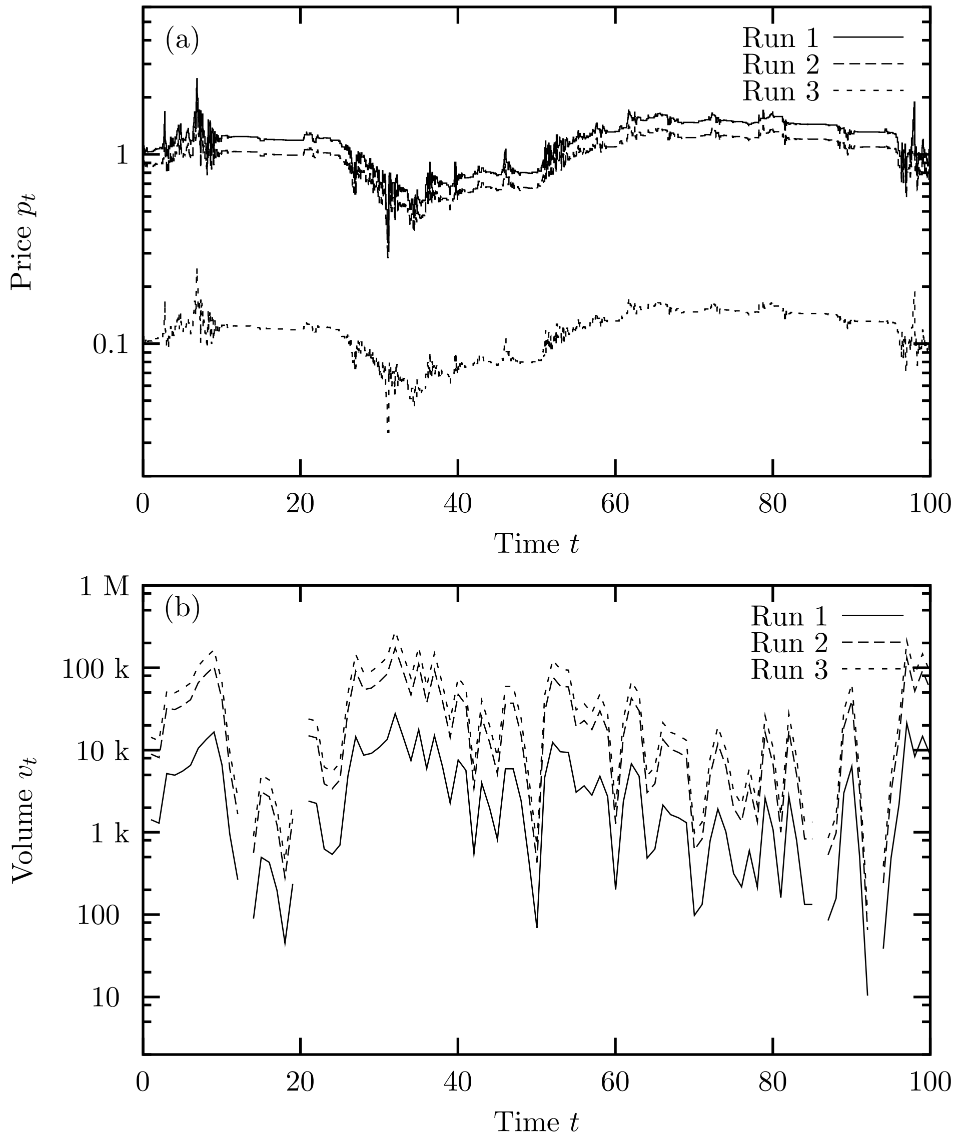

Figure 3.4: Comparison of time evolutions of (a) price and (b) volume for Runs 1, 2 and 3 as defined in Table 3.2. The price scales as the ratio of cash to shares and the volume scales as the number of shares. (Run 2 is offset to improve readability. The gaps in (b) denote periods of zero volume.)

Thus, we again have a model where the roles of C and S are only to set the price and volume scales, which are irrelevant anyway. Without loss of generality we can arbitrarily fix $C = \$1,000,000$ and $S = 1,000,000$ thereby reducing the degrees of freedom by two.

3.5.3 Further scaling

Having fixed the total cash and shares it is possible to show that no further scaling arguments are possible—there is no way to transform the model parameters such that the dynamics are invariant. We begin by noticing that, subject to C and S being fixed, each agent's cash c and shares s must also be fixed—constant under any transformation.

Under transformation a trade occurring between times t_- and t_+

$$\begin{aligned} c(t_+) &= c(t_-) - p\Delta s && (3.59) \\ s(t_+) &= c(t_-) + \Delta s && (3.60) \end{aligned}$$

becomes

$$\begin{aligned} c(t_+)' &= c(t_-) - p'\Delta s' && (3.61) \\ s(t_+)' &= c(t_-) + \Delta s'. && (3.62) \end{aligned}$$

But requiring $c' = c$ and $s' = s$ for all times immediately imposes the restrictions $\Delta s' = \Delta s$ and

$$p' = p. \tag{3.63}$$

So price must also be a constant under the transformation.

In particular, an agent's ideal price, given by Eq. 3.14, and limit prices, given by Eqs. 3.20–3.21, must remain unscaled, which immediately implies $i' = i$ and

$$f' = f, \tag{3.64}$$

so there is no way to rescale f without impacting the dynamics.

The constraint that i be invariant necessarily means that the evidence e must also be. After a number of news releases η_j and price movements the evidence is

$$e(t) = r_n \sum_j \eta_j + r_p \log \frac{p(t)}{p(0)}. \tag{3.65}$$

Therefore e is only invariant if the number of news releases is the same (requiring $\tau_n' = \tau_n$) and

$$\begin{aligned} r_n' &= r_n && (3.66) \\ r_p' &= r_p. && (3.67) \end{aligned}$$

The conclusion which may be drawn is that there is no transformation of any model parameters which leaves the dynamics invariant.

3.6 Parameter tuning

In the absence of scaling arguments to reduce the parameter space further we must use tuning methods to choose appropriate parameter ranges. (It is acknowledged this weakens the model's results somewhat, but it is necessary in order to establish a sufficiently small parameter space for experimentation.)

3.6.1 News response

In this section we will explore the effect of varying the news response parameter r_n. Let primes denote values after the arrival of a news event and unprimed quantities the same values before its arrival. If we could neglect price response ($r_p = 0$) then the price fluctuations would behave as

$$\log \frac{p^{*\prime}}{p^*} = e' - e = r_n \eta \tag{3.68}$$

where η is the cumulative news in the interval. So we would expect the price to have a log-Brownian motion with a standard deviation of $r_n\sqrt{t}$ (because η has a variance t).

However, the agents' response to price movements clouds the picture somewhat. Accounting for both news and price response, the evidence changes as

$$e' - e = r_n \eta + r_p \log \frac{p'}{p}. \tag{3.69}$$

If we assume, for simplicity, that the agent begins and ends at its ideal price ($p = p^*$ and $p' = p^{*\prime}$) then the relationship becomes

$$\log \frac{p^{*\prime}}{p^*} = \frac{r_n}{1 - r_p}\eta. \tag{3.70}$$

Although the assumption is too restricting for the above equation to accurately describe price movements it at least sets a scale for the dependence of the price movements with respect to the model parameters.

To test Eq. 3.70 the price series of Run 1 is reproduced in Fig. 3.5 along with the expected price from the equation (initialized at $p(0) = 1$). The graph indicates a rough agreement between the series but with a systematic error for prices far from \$1, indicating the equation does not completely capture the dynamics.

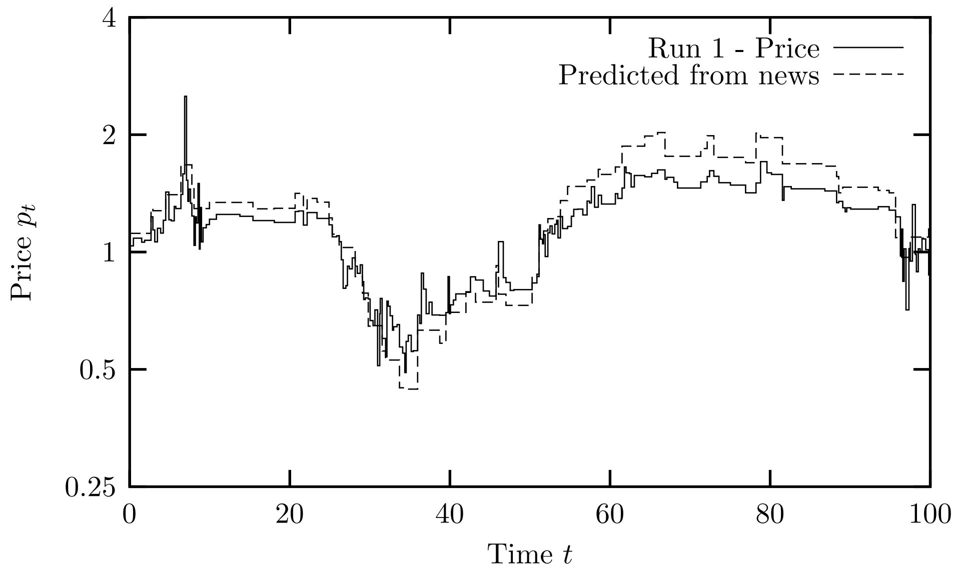

Figure 3.5: The price series generated by Run 1 is compared with the expected price generated by Eq. 3.70, showing rough agreement (though with systematic deviations).

Nevertheless, it is sufficient for the following purpose: it at least allows us to choose the scale of the news response r_n such that the fluctuations are on roughly the same scale as observed in real markets. If the log-return r (not to be confused with responsiveness) over one day obeys

$$r(t) \equiv \log \frac{p(t)}{p(t-1)} = \frac{r_n}{1-r_p}(\eta(t) - \eta(t-1)) \tag{3.71}$$

where $\eta(t)$ is the cumulative news up to time t, then the standard deviation σ_r of the returns is

$$\sigma_r = \frac{r_n}{1-r_p}\sqrt{\mathrm{Var}\,[\eta]} \tag{3.72}$$

where $\mathrm{Var}\,[\eta]$ is the news variance over one day, which was set in Section 3.3.2 to 1. Rearranging the relationship, we can set a scale for the news response

$$r_n = (1-r_p)\sigma_r. \tag{3.73}$$

Daily returns for the New York Stock Exchange over 26 years covering the period 1962–1988 [60] give $\sigma_r = 0.00959$ while the nine individual stocks studied in Chapter 6 are somewhat more variable with $\sigma_r = 0.036 \pm 0.014$ (see Table 6.4). A rough guideline, then, is $r_n = (1-r_p)\cdot 0.02$. Arbitrarily taking $\langle r_p\rangle \approx 1/2$ suggests $\langle r_n\rangle = 0.01$. For all future simulations a range of $r_n = 0.01 \pm 0.01$ will be used, unless otherwise stated.

3.6.2 Friction

In this section the effect of changing the friction parameter will be explored.

One of the effects of changing the friction has already been presented in Section 3.2.10 where it was found that, in order to preserve the number of trades per day, it was necessary to rescale time with Eq. 3.39. This has an appealing interpretation: as friction (or cost per trade) increases the trade rate drops.

However, this effect is rather trivial and, by rescaling time via Eq. 3.39, it can be ignored. Nevertheless, f still influences the dynamics in subtle ways.

To choose the friction we again appeal to real market structure. Most markets prefer trade quantities to be in *round lots* or multiples of one hundred shares. So the minimum trade in DSEM should consist of one hundred shares. Of course, the trade quantity depends on how many shares an agent holds and how the dynamics have unfolded. But, at the very least, we can impose this condition initially (at $t = 0$) because then we know each agent's portfolio and the market price precisely.

Assuming the total cash C and shares $S(= C)$ are distributed equally, each agent begins with an ideal investment fraction $i = 1/2$, ideal price $p^* = 1$, and limit prices $p_{BS} = (1+f)^{\pm 1}$. A trade will be initiated at one of the limit prices so the quantity of shares to be traded, from Eq. 3.19, is

$$\begin{aligned} s^* - s &= \frac{S}{2N}\left((1+f)^{\pm 1} - 1\right) && (3.74) \\ &\approx \pm\frac{Sf}{2N}. && (3.75) \end{aligned}$$

Imposing the round lot restriction $\Delta s_{min} = 100$ sets the friction at

$$f = 2\frac{\Delta s_{min} N}{S} = 200\frac{N}{S}. \tag{3.76}$$

From this argument it appears that the friction should increase with the number of agents N in the model. However, this is just an artifact of having fixed the total number of shares S. As N is increased each agent receives fewer shares since S is constant so it must wait for a larger price move before trading, so that it can trade a full lot. But fixing S in this way is somewhat unnatural. It is more natural to expect that if more agents are involved in a certain company then the company is probably larger and has more shares allocated. So S should probably have scaled with N.

But the effect of scaling S can be mimicked by scaling the round lot as $\Delta s_{min} \propto 1/N$. Then the friction is a constant, regardless of N. The most common system size to be used in this research will be $N = 100$ so, given $S = 1,000,000$, a good scale for the friction is $f = 0.02$. However, to incorporate heterogeneity future simulations will use $f = 0.02 \pm 0.01$.

3.6.3 News interval

Previously it was argued that the average news release interval τ_n should be on the order of a day. After more reflection the value of one day seems even more appropriate given the strong daily periodicity in real markets; daily market openings and closings, daily news sources (such as newspapers), and human behaviour patterns are just a few examples of daily cycles which influence market dynamics.

Another advantage of using an interval of one day in DSEM is that it strengthens the connection with CSEM, in which all events occur simultaneously once per day.

For these reasons the news release interval will be fixed at $\tau_n = 1$ (day).

3.6.4 Finalized parameter ranges

Via rescaling and tuning we have greatly narrowed the allowed ranges of the parameters. The final ranges which will be used in all further simulations are shown

Symbol	Interpretation	Range
	Market parameters	
N	number of agents	2+
C	total cash available	\$1,000,000
S	total shares available	1,000,000
τ_n	average interval between news releases	1
	Agent parameters	
f_j	friction of agent j	0.02 ± 0.01
$r_{n,j}$	news response of agent j	0.01 ± 0.01
$r_{p,j}$	price response of agent j	$\langle r_p \rangle \in (0, 1)$

Table 3.3: As Table 3.1 except with updated parameter ranges. These ranges will be used in subsequent simulations. All parameters except N and r_p are firm.

in Table 3.3. All the parameters except the number of agents N and the price response r_p have well-defined values (or ranges, in the case of agent-specific parameters). These two parameters will be the subject of further examination in the next chapter.