

**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

© Hendrik J. Blok, 2000

Chapter 2

Centralized Stock Exchange Model

2.1 Inspiration

In this chapter we will explore the *Centralized Stock Exchange Model* (CSEM), a microscopic model which is built upon the premise of *centralization*; each agent on the market is restricted to trading with a single, monopolistic *market maker* who has complete control over the execution price. No direct trades between agents are allowed. This situation approximates some actual, thinly traded stocks on the New York Stock Exchange (NYSE) and other markets [26].

There are two reasons this approach was chosen: firstly, there exists a significant collection of literature following this methodology [27–36]. I hoped to familiarize myself with this literature by constructing a model along the same vein.

Secondly, it allows for the construction of very simple agents. By having the trading price set exogenously the agents need only react rather than formulate their own trading schedules. In particular, the standard game theoretic approach is applicable only to reactive agents, as will be seen.

Hence, the development of CSEM was a natural starting point for my research.

2.2 Theory

In this section the model's structure will be explained.

2.2.1 Assumptions

I will begin, for the sake of clarity, by laying out some common assumptions used in CSEM.

Heterogenous agents

The market consists of many agents interested in trading. If all the agents had identical beliefs then we might expect their actions to be identical. Hence, we would effectively have a market of just one meta-agent unable to execute any orders. Similarly, if any subset of the population is homogeneous then that subset can be equally well represented by a single agent.

Therefore, it is natural to require that all the agents be unique. Notice that heterogeneity can arise from imperfect rationality or incomplete knowledge, qualities which seems reasonable for the simple agents which will be constructed. In most cases the agents will differ in fundamental parameters describing their preferences but transient differences alone (such as cash or shares held) may be allowed too, provided these factors influence the agents' actions.

Single risky asset and single riskless asset

For simplicity a market consisting of just one risky asset (public company stock) and one riskless asset (cash, for instance) will be used.

The total number of shares available on the market will be conserved. Since the company pays no dividend the stock has no fundamental value and stock price is maintained solely by expectations of satisfactory returns on the sale of shares. (The stock must at least have a chance of paying a dividend eventually or the stock price will be identically zero for all time, but the payout date is assumed to be far in the future.)

For simplicity, the stock price will be assumed to be a continuous variable. (In contrast, real stock prices are discretized, but on a sliding scale—dollar stocks usually have increments of one sixteenth of a dollar but penny stocks may be incremented by one tenth of a penny.)

The riskless asset (which we will call *cash*, though it could as easily represent some other stable equity such as gold) is defined to have a fixed intrinsic value in terms of which the value of the stock is measured. (By measuring all value in terms of cash some of the difficulties of utility theory in comparing utilities of disparate objects [37] are sidestepped.)

The total cash in the market will also be conserved. To achieve this, cash will pay no interest and no commissions will be charged on trades. This restriction

may be unrealistic but it has a significant advantage: the ratio of cash-to-shares is conserved. This means that the market can avoid moving into a regime dominated by one or the other and instead establish a balance between the two. For instance, if transaction costs were implemented cash would flow out of the system and, eventually, most of each agent's wealth would be held in stock. Conversely, if interest was paid on cash, eventually the market might be cash-dominated. In either case it is conceivable that the dynamics of the market would change, adding a complicating factor. Fixing the amount of cash in the system simplifies the model and allows for the collection of large datasets.

No intraday trading

CSEM uses a trading model which assumes all trades are executed only once daily (simultaneously). This approach is common in the literature [27–30, 33, 35] and mimics trading which occurs in real markets on unprocessed orders before opening each day.

Centralized trading

As mentioned above, the agents in this model are restricted to trading only with a single, monopolistic market maker or specialist. They are not allowed to trade directly with each other. This has an empirical basis but is also a simplifying factor. A discussion of how trading is implemented follows.

2.2.2 Utility theory

Each agent can adjust a portfolio consisting of s shares of a single risky stock and riskless cash c . If the share price is p then the agent's total capital at time t is $w_t = c_t + p_t s_t$. With interest and fluctuations in the stock price the agent's capital after one day (defined as one time unit) becomes

$$w_{t+1} = c_t + p_{t+1} s_t \tag{2.1}$$

$$= w_t + [p_{t+1} - p_t] s_t \tag{2.2}$$

and the trading behaviour reduces to an optimization problem with respect to the holdings s_t .

If the agent could know what tomorrow's price of the stock p_{t+1} will be in advance, finding the optimum strategy would be trivial: if $p_{t+1} > p_t$ then move all one's capital into the stock, otherwise move it all into cash. But of course the future price is unknown. Nevertheless each agent assumes it is a stochastic variable and

has some expectations of the underlying probability distribution, based on historical prices.

A naive goal, then, might be to maximize one’s expected future wealth $\langle w_{t+1} \rangle$ with respect to one’s current holdings s_t . Unfortunately, Eq. 2.2 simply tells us to invest all our capital into stock if $\langle p_{t+1} \rangle > p_t$ and otherwise into cash, almost exactly as before. The problem with this approach is that it doesn’t factor in *risk*. What if there was a non-zero probability that the stock price would crash $\Pr(p_{t+1} = 0) > 0$? Then, under repeatedly application of this strategy the agent would eventually lose all its wealth with certainty. Even if the price can’t drop to zero (which it can’t if there is any expectation of a non-zero price in the future) this strategy can perform poorly, particularly if the price is a multiplicative stochastic process [38] because it assigns disproportionate weights to extremely unlikely events which would have exorbitant payoffs. This strategy is said to be *risk neutral*.

We define our agents as simple expected utility maximizers where the *utility function* is monotonically increasing with wealth but has a negative second derivative (concave)

$$\frac{dU}{dw} > 0, \quad \frac{d^2U}{dw^2} < 0. \quad (2.3)$$

These requirements for a utility function are well established within financial economics [37, 39] and basically mean that an agent is unwilling to make a “double-or-nothing” wager of any amount if the odds are even. (Notice that the risk neutral agent $U = w$ would be ambivalent towards this wager and a *risk preferring* agent $\frac{d^2U}{dw^2} > 0$ would willingly take the wager.)

Exponential utility function

An often-chosen form [19] is the exponential utility $U(x) = -e^{-\alpha x}$ or equivalently (because utilities are defined only up to a linear transformation [37])

$$U(w) = w_{goal} \left(1 - e^{-w/w_{goal}} \right) \quad (2.4)$$

where w_{goal} is called the goal wealth and sets a natural scale for the utility. As shown in Fig. 2.1, the utility crosses over from a linear dependence on w at small wealth $U(w \ll w_{goal}) \approx w$ to an asymptote at large wealth $U(w \gg w_{goal}) \rightarrow w_{goal}$. The interpretation of w_{goal} as a “goal wealth” is justified because below w_{goal} the agent is willing to take risks for the chance of high payoffs but above w_{goal} it sees little reward in amassing greater wealth, being more concerned with maintaining its current level.




Figure 2.1: The exponential utility function defined in Eq. 2.4 is often applied in finance. The goal wealth parameter w_{goal} implicitly sets the risk aversion.

2.2.3 Optimal holdings

The exponential utility function is useful because it provides an analytic solution to the maximization problem [16] if we assume tomorrow's wealth w_{t+1} is Gaussian distributed (a reasonable assumption by the Central Limit Theorem, if it is a cumulation of many additive stochastic components). Then the expectation of the future utility is

$$\langle U(w_{t+1}) \rangle = \int dw_{t+1} U(w_{t+1}) \Pr(w_{t+1}) \quad (2.5)$$

$$= w_{goal} \left[1 - \exp \left(\frac{\text{Var}[w_{t+1}]}{2w_{goal}^2} - \frac{\langle w_{t+1} \rangle}{w_{goal}} \right) \right] \quad (2.6)$$

which is maximized by simply minimizing the argument of the exponential.

The future wealth depends on the price movement through Eq. 2.2 so the mean and variance become

$$\langle w_{t+1} \rangle = w_t + s_t \{ \langle p_{t+1} \rangle - p_t \} \quad (2.7)$$

$$\text{Var}[w_{t+1}] = s_t^2 \text{Var}[p_{t+1}]. \quad (2.8)$$

Eq. 2.6 can be maximized with respect to the free variable s_t yielding the

optimum quantity of shares to hold,

$$s_t^*(p_t) = \frac{w_{goal}(\langle p_{t+1} \rangle - p_t)}{\text{Var}[p_{t+1}]} \quad (2.9)$$

with the additional constraints $s_t^* \geq 0$ (no short selling) and $w_t \geq p_t s_t^*$ (no borrowing cash). The agent's strategy is to sell shares if $s_t > s_t^*$ or buy if $s_t < s_t^*$. The above equation is intuitively appealing: only hold shares if the expected return on your investment is positive and decrease your investment when the uncertainty (variance) is large (indicating an aversion to risk).

2.2.4 Risk aversion

The goal wealth w_{goal} in the utility function sets an undesirable, arbitrary scale for the agents behaviour: they will become increasingly risk neutral as their wealth falls far below this scale, and conversely, increasingly risk averse far above it. The arbitrary scale can be removed by setting the goal wealth proportional to the current wealth

$$w_{goal} = \frac{w_t}{a} \quad (2.10)$$

where a is a dimensionless constant which describes risk aversion (which increases monotonically with a).

Notice that introducing the dependence on the current wealth does not interfere with the optimization problem because w_t is a constant at any time t , independent of any changes in the portfolio s_t (assuming no trading costs). Therefore the optimal portfolio simply becomes

$$s_t^*(p_t) = \frac{w_t(\langle p_{t+1} \rangle - p_t)}{a \text{Var}[p_{t+1}]} \quad (2.11)$$

From Fig. 2.1 it is clear that the extremes of intense risk aversion and risk neutrality can be avoided by choosing a on the order of unity. A rough estimate provides an even more precise scale: empirically, the market appears to prefer to divide wealth equally between cash and stock when the annual expected return is 8% better than cash with an uncertainty on the order of 25%:

$$\langle p_{t+1} \rangle \approx (1 + 8\%)p_t \quad (2.12)$$

$$\text{Var}[p_{t+1}] \approx (25\%p_t)^2 \quad (2.13)$$

$$\Rightarrow s_t^* \approx \frac{1}{2} \frac{w_t}{p_t} \quad (2.14)$$

where t is scaled by years instead of days (but this does not interfere with the argument). The a -value to satisfy these conditions is $a \approx 2.5$.

Thus, the first agent-specific parameter introduced in CSEM is the risk aversion a which is constrained to lie in $a \in [1, 3]$.

2.2.5 Optimal investment fraction

For ease of comparison with the Decentralized model (to be presented in Chapter 3) the above discussion will be presented in terms of the fraction of one's wealth invested in stock. The investment fraction i_t at time t is given by

$$i_t = \frac{s_t p_t}{w_t} \quad (2.15)$$

and the *optimal* investment fraction is denoted by i_t^* .

Let us also define the return on investment from time t to $t + 1$:

$$r_{t+1} = \frac{p_{t+1} - p_t}{p_t} \quad (2.16)$$

which has a mean and variance (given a known current price p_t)

$$\langle r_{t+1} \rangle = \frac{\langle p_{t+1} \rangle - p_t}{p_t} \quad (2.17)$$

$$\text{Var}[r_{t+1}] = \frac{\text{Var}[p_{t+1}]}{p_t^2}. \quad (2.18)$$

Then, substituting Eq. 2.11 into Eq. 2.15, we find that the optimal investment fraction is

$$i_t^* = \frac{\langle r_{t+1} \rangle}{a \text{Var}[r_{t+1}]} \quad (2.19)$$

with the constraints $0 \leq i_t^* \leq 1$.

This relation has some intuitively attractive properties:

1. All else being equal, given two agents with different risk aversions, the one with the higher aversion will invest less.
2. Only invest if the expected return is strictly positive, and invest in proportion to it.
3. As your certainty of a good return increases (variance decreases), increase your investment.

However, it also has one glaring fault: when the expected return exceeds some limit,

$$\langle r_{t+1} \rangle \geq a \text{Var}[r_{t+1}] \quad (2.20)$$

the recommendation is to invest *all* capital in the stock, despite risk. This arises because the agents assume the returns are Gaussian-distributed, with no higher moments than the variance, but as we will see, higher moments do exist, increasing the risk.

Investment limit

To avoid complications of this kind a limit of δ is imposed: the investment fraction is constrained to lie within $i \in [\delta, 1 - \delta]$. Hence agents never take an absolute stance of investing all their money or withdrawing it all from the market. The motivation for this restriction is purely mathematical: it prevents the occurrence of all the agents simultaneously selling all their stock and driving the price down to zero (or conversely, selling all their cash and driving the stock price up to infinity).

As the population size increases, the probabilities of these events diminish simply as a result of fluctuations so the parameter δ becomes less important. To minimize its impact on the dynamics it should be assigned a small, positive value. Mathematically, however, it is allowed to be as large as 1/2 in which case the investment fraction would be a constant 1/2, never responding to Eq. 2.19.

Thus, the second agent-specific parameter is δ which is constrained to lie within $\delta \in (0, 0.5)$.

2.2.6 Forecasting

With Eq. 2.19 the optimization problem becomes one of forecasting one's future return r_{t+1} . In order to solve the optimization problem estimates of the expectation and variance of one's future return are required. The only information available to the agents is the history of returns so a reasonable choice is to try and extrapolate the series forward in time.

Although more complicated forecasting algorithms involving nonlinearity and chaos exist [40–45], I chose to extrapolate a simple curve-fitting algorithm to produce forecasts. The goal of this model is not to test complicated forecasting models but to understand the effect of interactions between many simple investors, so the forecasting algorithm need only be adequate, not optimal. Linear least-squares curve fitting is well understood so we don't have to worry about it generating unexpected side-effects in the dynamics.

The time series could be represented by a few parameters, one being the raw prices. However, a natural choice is the returns (as defined by Eq. 2.16) because a Gaussian-distributed future wealth w_{t+1} was assumed. This assumption can be validated by assuming a Gaussian distribution for returns as well, because Eq. 2.2 can be written as

$$w_{t+1} = (1 - i_t)w_t + i_t w_t r_{t+1} \tag{2.21}$$

where the stochastic variable is the return r_{t+1} . Since least-squares fitting assumes Gaussian errors, the returns are a convenient choice. Note that assuming a Gaussian

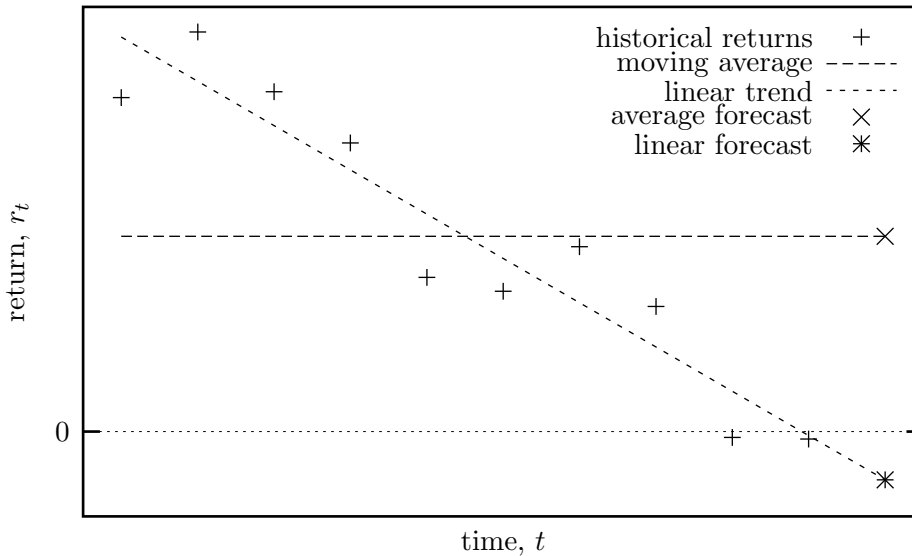


Figure 2.2: Demonstration of forecasting via polynomial curve extrapolation. Shown are forecasts produced by a simple moving average and a linear trend. The linear trend is able to anticipate reversals in returns.

distribution of returns is equivalent to a log-Brownian price series as is observed empirically on long timescales [46] (with interesting deviations on short timescales).

For simplicity, only low-degree polynomials will be used as fitting functions. The degree zero polynomial, a simple moving-window average, is already robust enough to project exponential growth in the stock’s price. Increasing to degree one (linear) also gives the agents the ability to forecast trend reversals (such as an imminent crash, as shown in Fig. 2.2) *assuming* the return history has meaningful trends.

By choosing higher degree polynomials we can effectively make the agents *smarter* (better able to detect trends in the return series) but, in practice, it is unreasonable to go beyond a degree two, quadratic fit. If too high a degree is chosen agents begin to “see” trends where none exist by fitting curves to noise.

Thus, the third agent-specific parameter, degree of fit d , is constrained to the integer values $d = 0, \dots, 2$.

Memory

Obviously, as time progresses and the latest returns are acquired, the older data in the time series become irrelevant. The standard methodology for handling this is to set a finite moving-window which only keeps the M most recent data points, discarding the rest. Then the curve fitting is performed only with respect to the remaining data. However, this technique has a drawback: it suffers from shocks as outliers (strongly atypical data) get dropped from memory.

To minimize this effect I constructed a method which uses an exponentially decaying window rather than the square window described above. The contribution of each point to the curve fit is weighted exponentially by how old it is. The technique is described in detail in Appendix A but a few points will be mentioned here:

The exponential weighting is characterized by a single parameter, the memory M (denoted by N^* in Appendix A) indicating the effective number of data points stored, which is approximately the decay constant of the exponential.

Using the exponential window allows compression of the data into just a few numbers regardless of the memory M and, as such, is computationally efficient in terms of storage and speed.

An agent's memory also says something about its expectations. A short memory produces fast responses to changes in returns and hence, more active trading. Conversely, a long memory results in slow variations in expectations and, therefore, slow changes in investment strategy. Hence, the memory implicitly also sets the (future) timescale, or horizon, over which the agent expects to collect.

As with standard curve-fitting the parameter M is required to be greater than the number of parameters to be fit ($= d + 1$ where $d \leq 2$ so $M \geq 10$ (two trading weeks) is satisfactory) but there is no maximum value. But to draw parallels with real markets it is reasonable to choose scales on the order of real market investors. Many online stock-tracking sites allow one to compare a stock's current value to its moving average over windows up to 200 trading days (almost one year).

Thus, the fourth agent parameter in CSEM is the memory M which is allowed to take on values in the range $M \in [10, 200]$ (between two weeks and roughly one year).

2.2.7 Fluctuations

To this point we have not explicitly identified the source of stochasticity. (Thus, since the simulation begins with no memory of any fluctuations no trading will occur whatsoever.) To mimic the noisy speculation which drives movements in real markets, stochastic fluctuations are introduced into CSEM. The fluctuations are

meant to represent the agents' imperfect information which can produce errors in their expectations of tomorrow's price. Given that the return-on-investment time series is already assumed to have Gaussian distributed errors a natural extension is to introduce normally-distributed fluctuations into the agents' forecasts

$$\langle r_{t+1} \rangle_\epsilon \equiv \langle r_{t+1} \rangle + \epsilon_t \quad (2.22)$$

where ϵ_t is a Gaussian-distributed stochastic variable with mean zero and variance σ_ϵ^2 .

It is assumed that agents are aware that their forecasts contain uncertainties so the variance of their forecasts is increased by

$$\text{Var} [r_{t+1}]_\epsilon \equiv \text{Var} [r_{t+1}] + \sigma_\epsilon^2 \quad (2.23)$$

since the forecasted return r_{t+1} is also assumed to be Gaussian-distributed (and the variance of the sum of two normally-distributed numbers is the sum of their variances).

Fluctuations are handled by determining a random deviate for each agent at each time step and adding it to the expected return, as discussed above. Once the deviate is chosen, it is a constant (but unknown by the agent) for that time interval, so the expected return is also constant. This is necessary for technical reasons (it keeps the agents' demand curves consistent for the auctioning process which will be discussed in Section 2.2.9) but it also seems intuitively reasonable—one would not expect an investor to forecast a different return every time she was asked (in the absence of new information).

The dynamics are driven solely by the presence of noise (as will be discussed below) so we require strictly non-zero standard deviations. On the other hand, the standard deviation also sets the typical scale of errors in the forecasted return. From personal experience, on a daily basis one would expect this error to be on the order of two percent. However, to fully explore the effect of the noise parameter CSEM will allow errors as large as 1/2 (which represents daily price movements up to $\pm 50\%$).

Thus, the fifth agent parameter introduced into CSEM is the scale of the uncertainty σ_ϵ which is chosen to lie within $\sigma_\epsilon \in (0, 0.5)$.

2.2.8 Initialization

The discussion so far has focused on how the agents evolve from day to day. But we must also consider in what state they will be started. It is important to choose starting conditions which have a minimal impact on the dynamics or a long initial transient will be required before the long-run behaviour emerges.

The simulations will be initialized with N agents; each agent will have a fraction of some total cash C and total shares S available. The effect of different initial distributions of cash and shares will be explored, but—unless otherwise specified—the cash and shares will usually be distributed uniformly amongst the agents. This allows the simulations to test the performance of other parameters; that is, to see if there is a correlation between parameter values and income.

As mentioned above, agents will also be initialized to have zero expectation $\langle r_1 \rangle = 0$ and zero variance $\text{Var}[r_1] = 0$ of tomorrow's return-on-investment. However, this is subject to Eqs. 2.22–2.23 so the actual initial expectation is a Gaussian deviate with mean zero and variance σ_ϵ^2 .

The first trading day is unique in that there exists no prior price from which to calculate a return-on-investment (for future forecasts). So the first day is not included in the agents' histories. Thus, the dynamics for the first two days of trading are due solely to fluctuations.

In this section three market parameters were introduced: the number of agents N on the market, as well as the total cash C and total shares S which are initially divided equally among the agents (unless otherwise stated).

2.2.9 Market clearing

Having discussed how the agents respond to prices and choose orders we now turn our attention to how the trading price is set. As mentioned before, this model is *centralized* in the sense that the agents are not allowed to trade directly with each other but all transactions must be processed through a *specialist* or *market maker* [27, 28, 30–36].

In real markets, the role of the market maker is more complex than in this simulation: here the market maker simply negotiates a price such that the market *clears*; that is, all buyers find sellers and no orders are left open. (All mechanisms by which the market maker may make a profit have been removed from the simulation for the sake of simplicity.)

A simple way for the market maker to establish a trading price is via an auction process: repeatedly call out prices and receive orders until buy and sell orders are balanced. If buy orders dominate, raise the price in order to encourage sellers, and vice versa.

However, CSEM provides a simpler (and faster) method for arriving at the trading price. Assuming the market maker knows each agent is using a *fixed investment strategy* as given by Eq. 2.19, it can be deduced that the optimal holdings for

agent j (with cash c_j and shares s_j) at price p is

$$s_j^* = \frac{c_j + s_j p}{p} i_j^*. \quad (2.24)$$

Effectively, by reporting their ideal investment fractions i_j^* (and current portfolios (c_j, s_j)), the agents submit an entire demand curve (demand versus price) for all prices instead of just replying to a single price called out by the auctioneer.

The market maker's goal of balancing supply with demand can be achieved by choosing a price which preserves the total number of shares held by the investors:

$$0 = \sum_j (s_j^* - s_j) \quad (2.25)$$

$$= \frac{1}{p} \sum_j c_j i_j^* + \sum_j (i_j^* - 1) s_j \quad (2.26)$$

which has a solution

$$p = \frac{\sum_j i_j^* c_j}{\sum_j (1 - i_j^*) s_j} \quad (2.27)$$

where, the values i_j^* , c_j , and s_j are all from before any trading occurs on the current day.

So, instead of requiring an auction, the trading price is arrived at with a single analytic calculation. Note that this method is possible because the optimal investment fraction i_j^* does not depend on the current day's price but only on the history of prior returns. (Once the trading price is established, the latest price is included in the history and contributes to the determination of tomorrow's optimal investment fraction.)

Initial trading price

In general, the calculation of the trading price is complicated and depends intricately on the history of the run but there is a special case where it is possible to determine explicitly the expected trading price—the first day. Let us assume that the initial distribution of cash and shares is such that each agent has equal numbers of both so that Eq. 2.27 reduces to

$$p_0 = \frac{\sum_j i_j^*}{\sum_j 1 - i_j^*} \quad (2.28)$$

$$= \frac{\langle i^* \rangle}{1 - \langle i^* \rangle}. \quad (2.29)$$

To calculate the expected investment fraction recall that initially the return history is empty so the expected returns are simply Gaussian-distributed with mean zero and variance σ_ϵ^2 so, from Eq. 2.19,

$$i_j^* = \frac{\epsilon_j}{a\sigma_\epsilon^2} \equiv \frac{x_j}{k}, \quad (2.30)$$

defining $x = \epsilon/\sigma_\epsilon$, $k = a\sigma_\epsilon$ and assuming the risk aversion a and forecast uncertainty σ_ϵ are identical for all agents.

Neglecting the limits $i \in [\delta, 1-\delta]$ on the investment fraction ($\delta = 0$) simplifies the calculation of the expected investment fraction:

$$\langle i^* \rangle = \int_0^{i(x)=1} i(x) \Pr(x) dx + \int_{i(x)=1}^\infty \Pr(x) dx \quad (2.31)$$

$$= \frac{1}{\sqrt{2\pi}} \left[\frac{1}{k} \int_0^k x e^{-x^2/2} dx + \int_k^\infty e^{-x^2/2} dx \right] \quad (2.32)$$

$$= \frac{1}{\sqrt{2\pi}k} \left(1 - e^{-k^2/2} \right) + \frac{1}{2} \left(1 - \text{erf}(k/\sqrt{2}) \right) \quad (2.33)$$

where $\text{erf}(\cdot)$ is the error function.

Substituting this equation into Eq. 2.29 gives the trading price as a function of the single parameter $k = a\sigma_\epsilon$, as shown in Fig. 2.3. Notice the value of the stock drops with increased risk aversion or uncertainty of return, properly capturing the essence of risk aversion.

It is interesting to note that the price drops to zero as $p_0 \propto k^{-1}$ for large k . To see how this occurs, notice that as the parameter k approaches infinity the second term in Eq. 2.33 drops out (falling off faster than $1/k$), as does the exponential in the first term, leaving only

$$\langle i^* \rangle (k \rightarrow \infty) \approx \frac{1}{\sqrt{2\pi}k}, \quad (2.34)$$

which diminishes to zero rapidly. The power law tail in the price emerges from simply substituting this relation into Eq. 2.29.

Now we briefly review the structure of the model.

2.2.10 Review

The Centralized Stock Exchange Model (CSEM) consists of a number N of agents which trade once daily (simultaneously) with a single market maker, whose goal is to set the stock price such that the market clears (no orders are pending). In this section, the structure of the model will be reviewed.




Figure 2.3: The expected initial trading price depends only on the risk aversion multiplied by the uncertainty of returns, $a\sigma_\epsilon$. As the aversion or uncertainty increases the initial value of the stock drops.

The agents are simple utility maximizers which extrapolate a fitted polynomial to the return history to predict future returns and, therefrom, optimal transactions. Each agent has a portfolio of cash c and shares s and is characterized by the parameters listed in Table 2.1.

Algorithm

Events are separated into days. After the model has been initialized the agents place orders and have them filled once each day. The basic algorithm follows:

1. Initialization. Cash and shares distributed amongst agents. Agents clear histories.
2. Start of new day. Agents forecast return-on-investment from history (and noise).
3. Agents calculate optimal investment fraction and submit trading schedules (optimal holdings as a function of stock price).
4. Market maker finds market clearing price (supply balances demand).
5. Trades are executed.
6. Agents calculate stock's daily return-on-investment and append to history.
7. End of day. Return to step 2.

Parameters

For convenience all the variables used in CSEM are listed in Table 2.1. 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.

2.3 Implementation

The above theory completely characterizes CSEM. The model is too complex for complete analysis so it is simulated via computer. The model was encoded 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 from <http://rikblok.cjb.net/phd/csem/>.

Symbol	Interpretation	Range
Market parameters		
N	number of agents	2+
C	total cash available	
S	total shares available	
Market state variables		
p_t	stock price at time t	
v_t	trade volume (number of shares traded) at time t	
Agent parameters		
a_j	risk aversion of agent j	[1, 3]
δ_j	investment fraction limit of agent j	(0, 0.5)
d_j	degree of agent j 's fitting polynomial	0, 1, 2
M_j	memory of agent j 's fit	[10, 200]
$\sigma_{\epsilon,j}$	scale of uncertainty of agent j 's forecast	(0, 0.5)
Agent state variables		
c_j	cash held by agent j	
s_j	actual shares held by agent j	
s_j^*	optimum shares held by agent j	
$w_j(p)$	wealth of agent j at stock price p	
i_j	actual investment fraction of agent j	
i_j^*	optimum investment fraction of agent j	

Table 2.1: All parameters and variables used in the Centralized Stock Exchange Model (CSEM).

2.3.1 Pseudo-random numbers

Coding the model as it has been described is fairly straight-forward. The only complication is that modern computers are unable to produce truly random numbers (required for the fluctuations in the forecasts) because computers are inherently deterministic.

Many algorithms for generating numbers which *appear* random have emerged. A good pseudo-random number generator must have three qualities: it must be fast, it must pass statistical tests for randomness and it must have a long period. The period exists because there are only a finite number of states (typically 2^{32}) a random number may take on. Hence, it must eventually return to its original seed and once it does, since the series is deterministic, it is doomed to cycle endlessly. If the period is less than the number of times the generator is called within a single run, the periodicity will contaminate the dataset.

One of the earliest and simplest pseudo-random number generators is the *linear congruential generator* [20, Section 7.1] which is defined recursively for an integer I_j :

$$I_{j+1} = aI_j + c \pmod{m}. \quad (2.35)$$

While this algorithm is fast it is not a good choice because it exhibits correlations between successive values.

More complicated generators have been developed which pass all known statistical tests for randomness [20, 21]. One of these, the *Mersenne Twister* [22] is also fairly fast and has a remarkable period of $2^{19937} \approx 10^{6000}$. Unless otherwise specified, the Mersenne Twister will be the generator of choice for CSEM.

Seed

All pseudo-random number generators require an initial seed: a first number (I_0 in the linear congruential generator, for example) chosen by the user which uniquely specifies the entire set of pseudo-random numbers which will be generated. This seed should be chosen with care: using the same seed as a previous run will generate the exact same time series (all other parameters being equal).

CSEM is coded to optionally accept user-specified seeds or it defaults to using the current time (measured in seconds since midnight, January 1, 1970, GMT). Since no two simulations will be run simultaneously, this provides unique seeds for every run. Unless explicitly specified, the default (time) seed will be used in the simulations.

2.4 Parameter space exploration

With CSEM coded into the computer, time series data can be generated for numerical analysis. As presented CSEM requires at least eight parameters to fully describe it. To fully explore the space of all parameters, then, means exploring an eight dimensional manifold... a daunting task. Before starting any experiments, then, it would be a good idea to check if any of these dimensions can be eliminated.

2.4.1 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.

2.4.2 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

$$C' = AC \tag{2.36}$$

$$S' = BS. \tag{2.37}$$

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

To see how these rescalings affect the dynamics let us begin by assuming that each agent's ideal investment fraction i_t^* is unchanged (this will be justified below). Then from Eq. 2.27 the price is rescaled by

$$p'_t = \frac{A}{B} p_t \tag{2.38}$$

and each agent's total wealth is rescaled by

$$w'_t = Aw_t. \tag{2.39}$$

(The rescaling of price can be interpreted as the "Law of supply and demand" because when either cash or stock exists in overabundance, it is devalued relative to the rarer commodity.)

Thus, the optimal holdings become

$$s_t^{*'} = \frac{w'_t}{p'_t} i_t^* = B s_t^* \tag{2.40}$$

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
σ_ϵ	0.5	0.5	0.5
M	40 ± 20	40 ± 20	40 ± 20
a	2 ± 1	2 ± 1	2 ± 1
δ	0.001	0.001	0.001
d	1 ± 1	1 ± 1	1 ± 1
seed	-2	-2	-2

Table 2.2: Parameter values for CSEM Runs 1, 2 and 3.

and the volume an agent trades becomes

$$\Delta s'_t = |s_t^{*'} - s'_t| = B\Delta s_t. \quad (2.41)$$

To justify that the optimal investment fraction remains unchanged, recall that it depends only on the return series through Eq. 2.19. The return series, under rescaling, becomes

$$r'_t = \frac{p'_t - p'_{t-1}}{p'_{t-1}} = r_t \quad (2.42)$$

assuming the price series is rescaled by A/B . Thus, if the investment fraction remains unscaled then the price series is scaled by A/B , so the investment fraction remains unscaled. . .

This would be a circular argument except for the fact that the investment fraction is initialized by a Gaussian fluctuation, which depends only on the parameters a and σ_ϵ . Thus the investment fraction begins unchanged (under rescaling of C and S) and there exists no mechanism for changing it, so it remains unchanged throughout time.

So, when cash is rescaled by some factor A and shares by B , the only effects are:

1. Trading price is rescaled by A/B .
2. Trading volume is rescaled by B .

To clarify this point in the mind of the reader, three identical runs were performed, with the parameter values shown in Table 2.2. Notice that Run 2 is Run 1 repeated with the scaling factors $A = B = 10$, and Run 3 is Run 1 with




Figure 2.4: Comparison of time evolutions of (a) price and (b) volume for Runs 1, 2 and 3 as defined in Table 2.2. The price scales as the ratio of cash to shares and the volume scales as the number of shares. (In both plots Run 2 is offset to improve readability.)

$A = 1, B = 10$. The resulting time series, shown in Fig. 2.4, confirm the claim that price scales as A/B and volume scales as B .

Neither the absolute value of the price nor the volume are items of interest in this dissertation. Instead we are interested in fluctuations, in the form of price returns and relative change of volume. Neither of these properties are affected by rescaling the total cash or total shares so we are free to choose a convenient scale. I have arbitrarily chosen a market with $C = \$1,000,000$ total cash and $S = 1,000,000$ total shares, thereby reducing the degrees of freedom by two.

2.4.3 Investment fraction limit δ

The investment fraction limit parameter δ sets a bound on the minimum and maximum allowed investment fractions $\delta \leq i \leq 1 - \delta$. This is purely a mathematical kludge to prevent singularities which could otherwise occur in Eq. 2.27.

Effectively, δ sets an upper and lower bound on the price itself: assume the total cash and shares are equal ($C = S$). Then, the minimum price is realized when all agents want to discard their stocks, $i_j^* = \delta$ for all j , giving

$$p_{min} = \frac{\delta}{1 - \delta}. \quad (2.43)$$

Conversely, given maximal demand, $i_j^* = 1 - \delta$, the price will climb to a maximum of

$$p_{max} = \frac{1 - \delta}{\delta}. \quad (2.44)$$

So the choice of δ sets the price range for the stock. Obviously, to allow reasonable freedom of price movements the limit should be significantly less than one half, $\delta \ll 1/2$. To mimic the observed variability in some recent technology-sector stocks, a limit of $\delta = 0.001$ will generally be used, allowing up to a thousand-fold increase in stock value—except in Chapter 4 where we explore the effect of varying this parameter.

2.4.4 Risk aversion a and forecast uncertainty σ_ϵ

One's intuition may lead one to suspect that the risk aversion factor a and the forecast uncertainty σ_ϵ are over-specified, and should be replaced by a single parameter $k = a\sigma_\epsilon$ as was done to calculate the initial trading price in Section 2.2.9. However, a closer inspection of Eq. 2.19 demonstrates this is not quite true. The optimal investment fraction is

$$i_t^* = \frac{\langle r_{t+1} \rangle + \epsilon_t}{a(\text{Var}[r_{t+1}] + \sigma_\epsilon^2)}. \quad (2.45)$$

Parameter	Run 4	Run 5
N	100	100
C	\$1,000,000	\$1,000,000
S	1,000,000	1,000,000
σ_ϵ	0.25	0.5
M	40 ± 20	40 ± 20
a	2	1
δ	0.001	0.001
d	1 ± 1	1 ± 1
seed	-2	-2

Table 2.3: Parameter values for CSEM Runs 4 and 5.

Since σ_ϵ is the only parameter to set a scale for the returns, in Eqs. 2.22–2.23, it is reasonable to expect the returns to scale linearly with σ_ϵ so renormalizing gives

$$i_t^* = \frac{1}{a\sigma_\epsilon} \left[\frac{\langle r_{t+1} \rangle / \sigma_\epsilon + \epsilon_t / \sigma_\epsilon}{\text{Var}[r_{t+1}] / \sigma_\epsilon^2 + 1} \right] \quad (2.46)$$

where the second factor is invariant under rescaling of σ_ϵ .

Then, since a and σ_ϵ occur nowhere else in the model, one may expect that the simultaneous rescaling

$$a' = Ca \quad (2.47)$$

$$\sigma'_\epsilon = \sigma_\epsilon / C \quad (2.48)$$

would preserve the dynamics.

However, as the price series of Runs 4 and 5 (see Table 2.3) show in Fig. 2.5, there are small deviations which grow with time until eventually the time series are markedly different.

To see why this occurs, let us consider a simple thought experiment: Consider a run with equal amounts of cash and shares ($C = S$) where the last trading price—for the sake of convenience—is $p_{t-1} = 1$. Now assume that on the next day all the agents have negative fluctuations in their forecasts which drive their optimal investment fractions to their lower limits $i_t^* = \delta$. Then, from Eq. 2.27, the day's stock price will be given by Eq. 2.43 and the return will be

$$r_t = \frac{p_t - 1}{1} = -\frac{1 - 2\delta}{1 - \delta} \quad (2.49)$$

which does not scale with σ_ϵ as was hypothesized in the derivation of Eq. 2.46.




Figure 2.5: Comparison of time evolutions of price for Runs 4 and 5 as defined in Table 2.3. The price is not perfectly invariant under rescalings which preserve the constant $a\sigma_\epsilon$.

Occasional events like the one described in the above thought experiment are responsible for the deviations seen in Fig. 2.5. However, apart from these rare deviations (which, neglecting trends in the return history, should occur with decreasing frequency $1/2^N$ as the number of investors increases) the risk aversion parameter a and uncertainty σ_ϵ appear to be over-specified. Therefore, the risk aversions will always be chosen from a uniform deviate in the range $a \in [1, 3]$ and only the forecast error σ_ϵ will be manipulated—excepting the following section in which the relative performance of different values of a and σ_ϵ will be evaluated.

2.5 Parameter tuning

Thus far we have isolated three parameters (C , S , and δ) which can be fixed at particular values without loss of generality. We now want to choose reasonable ranges for the remaining parameters (σ_ϵ , M , a , and d). *Reasonable*, in this context, refers to agents with parameter combinations that tend to perform well (accumulate wealth) against dissimilar agents. These parameter combinations are of interest because one would expect that, in real markets, poorly performing investors who consistently lose money will not remain in the market for long.

Note that, as discussed in the Introduction, parameter tuning generally diminishes an explanatory model’s validity. This, however, does not quite apply in this case because we are not tuning the parameters in order to produce a model which better fits the empirical data (i.e., exhibits known market phenomena, such as fat tails and clustered volatility)—rather, we are simply trying to select “better” investors. However, it must be acknowledged that this *may* concurrently tune the simulation towards realism.

Further, the point of this exercise is not to completely specify the model but merely to avoid wasteful parameter combinations which should be driven out of the system by selective (financial) pressures. In the model, “dumb” agents (with parameter combinations which tend to underperform) will lose capital and may eventually hold a negligible portion of C and S . Hence, these agents won’t contribute to the market dynamics and will simply be “dead weight”, consuming computer time and resources. Hopefully, at this point the reader agrees that it would be helpful to cull “dumb” agents by finding the more successful parameter ranges.

It may be discovered, in the course of this investigation, that some parameters are irrelevant; they may be take on a wide variety of values with little or no impact on the dynamics. In this case, these parameters may be assigned arbitrary ranges without loss of generality.

To determine successful values, a large parameter space should be explored.

Parameter	Run 6
N	400
C	\$1,000,000
S	1,000,000
σ_ϵ	0.25 ± 0.25
M	105 ± 95
a	1.5 ± 1.5
δ	0.001
d	1 ± 1
seed	random

Table 2.4: Parameter values for CSEM Run 6.




Figure 2.6: Price history generated by CSEM with parameters listed in Table 2.4 (Run 6). The price almost reaches its theoretical maximum of \$999 (see Eq. 2.44) before collapsing. The agent state variables were sampled at the times indicated.

Parameter	Correlation with $\log w$			
	Sample 1	Sample 2	Sample 3	Sample 4
σ_ϵ	+	0	0	+
M	0	0	0	+
a	+	+	+	+
d	0	0	0	0

Table 2.5: Regression analysis of $\log w$ versus agent parameters for different samples of Run 6 (Table 2.4). The symbols indicate the sign of the regression-line slope, or zero if it is insignificant (relative to its standard error). The results indicate that a is positively correlated with wealth but σ_ϵ , M and d are largely irrelevant.

To this end a long data set was collected with more agents and with broader parameter ranges, as indicated in Table 2.4. The price history for the run is shown in Fig. 2.6.

The results were analyzed by looking for correlations between an agent’s wealth and the following parameters: forecast error σ_ϵ , memory M , risk aversion a , and degree of curve-fit d .

Note that the point of this work is not to determine an optimal investment strategy (set of optimal parameter values), but simply to establish reasonable ranges for these parameters such that the agents perform reasonably well. Thus, a complete correlation analysis is unnecessary. Instead, a simple graphical description of the results should be sufficient, with a simple regression analysis for emphasis.

Table 2.5 shows the results of linear regression analyses of $\log w$ versus agent parameters for different samples of Run 6. The logarithm of wealth is fitted to a straight line with respect to the parameter of interest and the sign of the slope is recorded. If the slope m has a standard error larger than 100% then the parameter is interpreted as being uncorrelated with performance. This method was constructed only because it lent itself to the computational tools available to the author. However, it is reasonable: recall that the linear correlation coefficient (which is typically used to test for correlations) is related to the slope $r \propto m$. Also, the standard error estimates the significance of the slope; a value greater than 100% suggests that the sign of the slope is uncertain.

The results of the analysis indicate that the risk aversion parameter a is positively correlated with performance (wealth). However, the forecasting parameters σ_ϵ (forecast error), M (memory) and d (degree of polynomial fit) appear to be uncorrelated with performance. Hence, these parameters can be set arbitrarily. The memory from Run 6 ($M = 105 \pm 95$) will be used in all further simulations to

Parameter	Run 7
N	400
C	\$1,000,000
S	1,000,000
σ_ϵ	0.025 ± 0.025
M	105 ± 95
a	1.5 ± 1.5
δ	0.001
d	1 ± 1
seed	random

Table 2.6: Parameter values for CSEM Run 7.

maintain diversity. However, the degree of the fitting polynomial will be constrained to $d = 0$ (a moving average) because it boosts simulation speed. The forecast error σ_ϵ requires further inspection.

Representative graphs of $\log w$ versus the parameters a and σ_ϵ (using Run 6: Sample 3) are shown in Fig. 2.7. The slopes are used to estimate correlations, as discussed above. The evidence suggests that risk aversion is positively correlated with performance. Hence, small values of a (high-risk behaviours) tend to underperform. Thus, the range of a is restricted to $a \in [1, 3]$ instead of $a \in [0, 3]$ as set in Run 6.

2.5.1 Forecast error

The only free parameter left is the forecast uncertainty σ_ϵ . Although Fig. 2.7 indicates no correlation between wealth and forecast error, a closer inspection reveals a small peak for the smallest errors $\sigma_\epsilon < 0.05$.

To test this range, a new dataset was collected with all the parameters as in Run 6 except the forecast error scaled down by a factor of ten, as indicated in Table 2.6. The time series, shown in Fig. 2.8, exhibits wildly chaotic fluctuations which regularly test the price limits (Eq. 2.44) imposed by δ . Since δ was an arbitrarily chosen parameter, we do not want it to significantly affect the dynamics as it does in Run 7.

Thus, the choice of $\sigma_\epsilon = 0.025 \pm 0.025$ causes problems. This issue will be revisited in Chapter 4.




Figure 2.7: Plot of agent wealth versus (a) risk aversion and (b) forecast error. The best fit lines have slopes 5.2 ± 1.4 (positive correlation) and 4.7 ± 8.8 (no correlation), respectively.




Figure 2.8: Price history generated by CSEM with parameters listed in Table 2.6 (Run 7). The series has the undesirable property that the price spends much of its history at or nearing its ceiling (\$999).

Symbol	Interpretation	Range
Market parameters		
N	number of agents	2+
C	total cash available	\$1,000,000
S	total shares available	1,000,000
Agent parameters		
a_j	risk aversion of agent j	2 ± 1
δ_j	investment fraction limit of agent j	0.001
d_j	degree of agent j 's fitting polynomial	0
M_j	memory of agent j 's fit	105 ± 95
$\sigma_{\epsilon,j}$	scale of uncertainty of agent j 's forecast	$[0, 0.5]$

Table 2.7: As Table 2.1 except with updated parameter ranges. These ranges will be used in subsequent simulations.

2.5.2 Finalized parameter ranges

The finalized ranges of the model parameters are shown in Table 2.7. The risk aversion and memory will always be assigned the shown ranges but the effect of varying N and σ_ϵ will be explored further in Chapter 4.

2.6 Discussion

In this section some observed properties of the model (both theoretical and empirical) will be discussed.

2.6.1 Fundamentalists versus noise traders

This model borrows heavily from other work in the area [27–36]. However, it differs from most of these papers in that it does not divide the traders into types. Many other models assign the agents one of two roles: either fundamentalists or noise traders [17, 32, 36, 47, 48]. Fundamentalists believe the stock has a real value (for instance, if it pays a dividend) and trade when they believe the stock is over- or under-valued. Noise traders (or *chartists*), on the other hand, have no interest in the stock’s fundamental value, but simply try to anticipate price fluctuations from the historical data, and trade accordingly.

CSEM deliberately eliminates the fundamental value of the stock so the agents are necessarily what would be called noise traders. Hence, the dynamics which emerge from the simulations are of a completely different nature than those mentioned above.

2.6.2 Forecasting

Table 2.5 indicates that the forecasting parameters are largely irrelevant to performance. This suggests that there are no serial correlations in the stock price and, therefore, no reward for increased effort to forecast (by increasing M and/or d). Whether the time series actually does have auto-correlations will be explored in Chapter 5. But the ineffectiveness of forecasting raises the question of whether a model based on forecasting is even relevant. Perhaps the agents would do better to ignore the return history and just rely on fluctuations to make their estimates of future returns. This may indeed be a valid argument but forecasting has another purpose—it adds a degree of heterogeneity to the agents through systematic differences in their investment fractions.

On the other hand, forecasting may provide a mechanism for herding. As the history develops, the returns may be correlated for short periods. If so, then

the agents may converge in opinion regarding future returns and act in unison, with significant consequences in the price history. For this reason, the forecasting algorithm will be retained.

2.6.3 Portfolios

Given an investment fraction i , wealth w , and stock price p , an agent's distribution of cash c and shares s is given by the relations

$$iw = sp \tag{2.50}$$

$$(1 - i)w = c. \tag{2.51}$$

Given the investment limits $\delta \leq i \leq 1 - \delta$, a linear relation between wealth and the maximum or minimum holdings of both cash and shares can be found. (Recall, agents are not allowed to sell all their shares or cash.) Fig. 2.9 shows the distribution of cash versus shares for the agents of Run 6: Sample 4. Notice that the agents almost exclusively hold extremal portfolios dominated by either cash or stock. Very few actually hold mixed portfolios. This indicates that Eq. 2.19, which gives an agent's optimal investment fraction, may be too sensitive. But the only freedom one has in reducing the sensitivity is through the parameters a and σ_ϵ , which have other consequences, as has been discussed.

2.6.4 Difficulties

Although this model showed promise, I had some technical and ideological problems which encouraged me to abandon it in favour of a different approach. On the technical side, as the reader can see, the number of parameters is somewhat unwieldy. Although some of the parameters could be determined, those remaining were difficult to manage. The investment fraction limit δ , for instance, is a necessary but unappealing result of the derivation, which imposes arbitrary limits on the stock price's range. Another difficult parameter is the forecast error σ_ϵ : if too large a value is chosen then the dynamics are dull and dominated by noise (see Fig. 2.5), but too small a value produces wildly chaotic behaviour completely unlike empirical market data (Fig. 2.8). This seems to be the critical parameter for determining the character of the dynamics, and the effect of varying it will be explored in more detail in Chapter 4.

One of the ideological problems was the use of parallel updating (all agents trading at a single moment each day). Evidence is mounting that employing a parallel updating scheme (without strong justification) introduces chaotic artifacts




Figure 2.9: Plot of agent wealth versus (a) cash and (b) shares held showing that most agents hold extreme portfolios of maximum cash and minimum shares, or vice versa. It appears that the method of calculating the investment fraction in CSEM (Eq. 2.19) is too sensitive to fluctuations. (It should be acknowledged the plots are truncated since the lowest wealth actually extends down to 10^{-25} , an unrealistic quantity since real money is really discretized with a minimum resolution of one penny.)

into the dynamics which are generally not observed in the actual, continuous-time system being modeled [14, 49–52].

Further, in this model the price of a share is artificially fixed by the market maker. In most markets the price is an emergent phenomena: auction-type orders are placed at hypothetical prices (eg. limit prices) and the price is realized when a trade occurs. Forcing the price to balance supply and demand destroys its emergent character.

For these reasons, this line of research was replaced with the model presented in the next chapter. Nevertheless, the Centralized Stock Exchange Model is included here because it follows a prevalent line of reasoning in stock market simulations and falls into many of the same pitfalls encountered by others [17, 27, 28, 30, 32, 34, 36, 48, 53]. Wherever possible, CSEM data will be analyzed alongside the output of the next model, the *Decentralized* Stock Exchange Model.