4. Financial market models with global interactions

4.1 A bottom-up approach

So far we have talked about stochastic processes, and have criticized the standard model of finance. The common theme of this criticism has been that standard finance theory has an Achilles heel, which is the limited applicability of the stochastic models typically employed to describe the underlying market movements. No matter how sophisticated one's model for portfolio management, derivatives pricing, risk analysis, etc., it is limited by the reliability and accuracy of the underlying model describing the market dynamics. It goes without saying that many modifications to these standard stochastic models have been suggested in the academic literature. How does one decide which modification will work best in a given market? Or how might we systematically create a new, stochastic supermodel? One way might be to make a detailed empirical comparison between the resulting statistics for various candidate stochastic models, and the stylized facts discussed in Chapter 3. However, this would necessarily involve a detailed comparison to past data for that particular market. The basic question then arises: will history actually repeat itself in this market?

A related problem arises if we are interested in understanding the evolution of a new market. By definition, there is now no history available. A recent example is that of the Euro prior to its launch—perhaps we would have liked to get a feel for the possible future evolution of its exchange rate against other major currencies. Is there any way to understand, at least in a general qualitative way, what will happen following such a launch? Alternatively, we might be regulators trying to work out what will happen when some global rules concerning the market we are monitoring or controlling, are changed. Again, we would like to get a feel for likely scenarios. An example of this would be the lifting of constraints on a given currency's exchange rate vs the US dollar—for example, the case of the Colombian Peso in the late 1990s. What will be the effect? Is there any way that we can say, at least in general terms, whether volatility is likely to increase, decrease, or stay the same? Maybe instead we are trying to work out what will happen following a particularly significant market event, or a given piece of bad news (e.g. the 2001–02 Enron scandal, or the 11 September tragedy in the United States). Is there any way we can foresee the likely extent of any resulting downturn, or the expected recovery time, or the volatility evolution in the immediate future?

These are all important practical questions, but obviously very hard to answer. Certainly no stochastic model of the type we have seen is ever going to answer these questions for us. After all, such stochastic models would require the model-builder to decide a priori which modifications are likely to be important in the future for that particular market. The model-builder is therefore forced to include his own forecast of how he believes the market will behave—in particular, he must make implicit assumptions about future trader activity. Going further, we can ask the following: no matter how sophisticated any stochastic supermodel that we develop, would we really be any closer to understanding how the markets are going to behave?

A science-motivated approach would suggest that we need instead to develop a microscopic understanding of what goes on in a market, in order to work out the best way to extend standard finance theory to address such questions. If it ain't broke don't fix it—but if it is broke, then fix it in a way that you understand. Looking to physics, we recall that thermodynamics—which studies the macroscopic behaviour of gases—lacked a firm foundation and understanding until the development of an atomistic theory based on statistical mechanics. So can we do the same as was done with the understanding of gases? In other words, can we hope to model the microscopic structure of markets based on how we think traders behave, and then use this to understand the observed macroscopic price dynamics? It seems like a hopeless task. But just as you do not need to know the detailed quantum properties of atoms in order to model the macroscopic average properties of a gas, so in the same way we would hope that one could develop a macroscopically realistic model of the market even though the microscopic details are grossly over-simplified. In short, maybe we can build market models which can be used to answer the above questions, but without having to include the characteristics of each individual trader. We do know that traders are a diverse bunch of people, with diverse opinions, time horizons, profit margins, trading capital, etc.—so maybe we can get away with a model which incorporates this diversity or heterogeneity, but without worrying about individual traders' preferences, likes, and dislikes. As we show in this chapter and Chapter 5, we can indeed get a significant distance towards reproducing the so-called stylized facts of Chapter 3 using very simple models of a population of heterogeneous agents (traders). Such models, in our opinion, open the way to next-generation finance. In addition, these models offer something significant in return to the scientific community: they are generic models of a complex system. As discussed in Chapter 1, the successful development of a theory to describe the dynamics of a complex system represents a fundamental outstanding problem within the sciences. For these reasons, we are going to spend some time discussing such models.

4.2 Two's company, but three's a crowd

It is clear that the fluctuations observed in financial time-series should, at some level, reflect the interactions, feedback, frustration, and adaptation of the markets' many and diverse participants. Using the analogy between markets and casinos, the trader becomes a gambler and the market becomes the game.[1] Since there are many market agents, this game is a multi-agent (multi-trader) game. The possible actions in this game become 'buy', 'sell', or 'do nothing'. The goal? This could be to win, or at least not to lose. In particular, this real-world market game will incorporate the following features.[2]

Limited, primarily global information. In a 2-player game we can usually work out what the decision of the other player was, just by looking at the pay-off we receive. However, in a game with 3 or more players, where we each receive just global information, it will not in general be possible to work out exactly what each agent did.[3] This leaves each player with an inherent uncertainty as to the strategies that each of the others used. In short, it is practically impossible for you to infer the microscopic strategies played by each of the other players, based only on global information concerning the winning outcome. In a market where there is no private exchange of information, the players see the price move—they then either win or lose based on their own decision to buy, sell, or do nothing, and the actual movement of the price. Any one trader can never work out exactly what another particular trader did.[4]

The IT infrastructure and communications in place in most of today's main financial centres is such that every large investor has access to essentially all the information that is available regarding any financial asset they are trading. In short, a good starting approximation is that all such information is public as opposed to private. This implies that all the agents base their investment decisions on the same information[5]: we refer to this information as 'global information' and assign it the variable $\mu[t]$.

[1] The interesting twist is that the spin of the roulette wheel, or deal of the cards, is actually determined by the aggregate action of the population of gamblers (agents) themselves.

[2] These features are consistent with the features of a complex system (see Chapter 1).

[3] If the players are competing, they are unlikely to phone each other to report on what they are doing. Even if they did, there would be no guarantees that such information was reliable.

[4] See note 3.

[5] In Chapter 5, we will look at the effect of imitation, which implies some degree of local information as to what other groups of agents are doing.

Many participants. Classical game theory tends to focus on games with $N = 2$ players. It then aims to deduce the equilibrium properties under which pay-offs are maximized. Since there are few players, and few strategies, there are a small number of such equilibria. It is therefore reasonable to expect that the players would recognize, and be able to evaluate, the details of such equilibria, and would act accordingly. However, when there are many players, this is not possible practically. For a simple N-player game, where each player has only $s = 2$ strategies, the pay-off diagram which must be analysed to deduce these equilibria has 2^N entries. For $N = 2$, it is just a 2×2 matrix and hence the equilibria are easy to compute: however, the complexity of this computation increases rapidly as N increases. No human player would be able to perform this computation for large N. It is therefore impossible for a financial market agent i to deduce his best investment strategy at any given time t (and hence his optimal order $a_i[t]$) without a complete knowledge of what orders $a_j[t]$ all the other agents $j \neq i$ will place. For this reason, agents will attempt to keep their orders secret from the other agents to avoid giving away any advantage. This has the result that the agents act independently of each other and must come to their own, inductive, conclusions about what their optimal order should be.

Although the market has many participants, it is not 'many' in the usual sense of statistical physics. For example, in a gas, there are of order 10^{23} atoms or molecules—however, there are probably of order only 10^2 market agents (traders) whose effect on the market is significant. It is these traders that we wish to model, since it is the aggregate of their actions which moves the price. The 10^2 traders in question include large banks, pension funds, institutions, and hedge funds.

Dynamical, as opposed to static. Financial markets are continually evolving, and one can in principle trade on very short timescales. Hence, a market represents a *repeated* game at the very least. The research on repeated N-player games that exists tends to focus on randomly picking out 2 players at each timestep and making them play a particular game. The system is then updated, a new pair of players chosen, and so on. Hence, there is no sense in which all N players are all playing at the same time, unlike in an actual market.

As stressed in Chapter 1, the issue of time and timing in markets is a crucial one. Time underlies the dynamics, introduces causality and hence produces a strict flow of events, that is, the action $\mathrm{A}[t_\mathrm{A}]$ cannot affect the action $\mathrm{B}[t_\mathrm{B}]$ if $t_\mathrm{A} > t_\mathrm{B}$. Even though individual agents may place their orders at slightly different times, the rate-limiting step is the market-maker. The market-maker can only work at a finite speed (even if the job is done electronically) and so will batch together all the orders which have been submitted since the order book was last cleared. This makes the system inherently discrete in time, and hence we can use the simplifying step

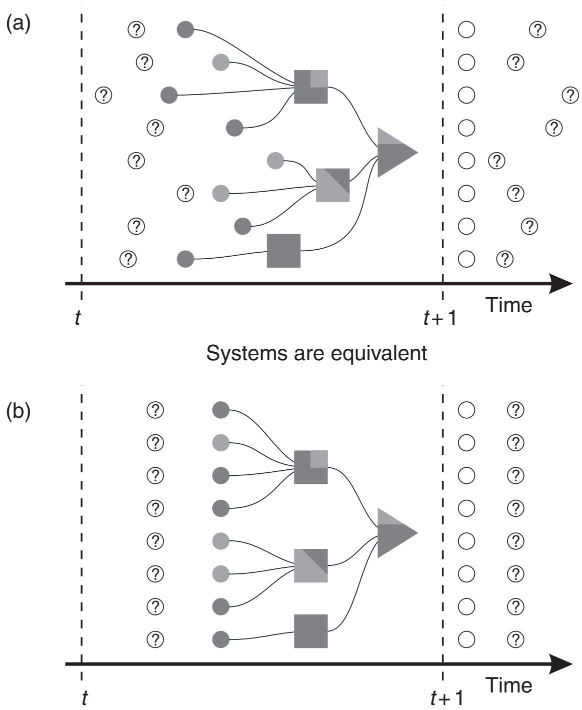

Fig. 4.1 Schematic diagram, with key as per Fig. 1.2, showing the equivalence of (a) the real market wherein orders can be placed by agents at different times, and (b) the simplified system wherein all agents place orders simultaneously.

of assuming that the agents all submit their orders at the same time. This assumption is *not* an approximation since the two systems are equivalent, as shown in the schematic diagram of Fig. 4.1. We do however propose making an approximation to the real market in the way that the orders are executed. Specifically, we would like to impose the tight temporal framework of decision–order–price update–order execution, shown in the schematic diagram of Fig. 1.2. The imposition of this framework is equivalent to asserting that all orders are market orders, since they are executed immediately after the asset price has been updated by the market-maker. This approximation simplifies the model framework greatly and is likely to be a reasonable approximation for high frequency trading where agents form an opinion as to which way the asset price will move over the next few trades. Consequently, the only thing that is important to these agents is that their trade be immediate in order to benefit from the movement. We will therefore only consider orders of the form $a_i[t, x[t+1]] = a_i[t]$.

Adaptation. A simple framework for an inductive agent is to allocate to that agent a number s of 'strategies' R, which map the available global information $\mu[t]$ at time t to an investment decision or 'action' $a_R^{\mu[t]}$. The provision of a set of s strategies allows the agents to inductively reason what will be the best order to make at time t, by looking at how well the strategies have performed previously. This also gives an agent scope for being able to adapt his behaviour to the current market conditions. A strategy that performed well at one stage in the past can become superseded by another as the nature of the market's behaviour changes.

Feedback. The global information $\mu[t]$, which contains all relevant information about the asset being traded, must contain variables such as the previous traded prices of the asset and the previous traded volumes. However, these variables are themselves generated by the system as output, so we have the situation that the system output is subsequently being used as part of the system input. This is feedback. If the approximation is made that the population of agents is dominated by chartists—who typically only consider patterns in previous asset prices and volumes to be relevant information—then the feedback in the system is perfect in the sense that there are no external factors influencing the agents' orders. The market therefore just feeds off itself. Not only is this a particularly interesting type of system to investigate, but also one of increasing relevance to the real financial markets. As discussed in Chapter 1, the more modern, liquid, and speculatory markets are precisely the ones that are becoming dominated by chartists.

Competition. One would expect the nature of the market to be one of outright competition. However, whether the agents might be cooperating with each other at some level, is a subtle point. The aim of each and every agent is to make a profit from their speculatory trading. However, it is not necessarily the case that profit earned by one agent will have the consequence of a direct loss for another; the system is not closed. This is mainly due to the freedom of the market-maker to manipulate the price of the asset at will. This process requires no flow of capital in the system and yet can generate or diminish the collective wealth of the population of agents with positions in the asset. Agents who are buying assets will tend to be pleased if the traded price for their order is lower than the asset price at the time they ordered. Similarly, agents who are selling will tend to be pleased if the traded price for their order is higher than the asset price at the time they ordered. However, an excess of buyers/sellers moves the asset price up/down, therefore the group of agents who 'wins' is the sellers/buyers as they get to sell/buy at a higher/lower price. At the point of trading, therefore, there will be competition to be in the minority group of 'winners'. However, the agents could also cooperate in order to manipulate the market-maker and so increase the value of their positions. We will be looking at these aspects related to how agents 'win', throughout this chapter.

4.3 'To bar, or not to bar . . .'

The economist Brian Arthur proposed the following problem[6] to embody why the economy, and in particular, a financial market, is so 'complex'. It incorporates the necessary ingredients mentioned in Section 4.2, which are characteristic of the financial market 'game', and was specifically designed to illustrate the shortcomings of standard economic theory. In particular, it shows the importance of out-of-equilibrium behaviour in a financial system, where agents behave inductively rather than deductively since they only have limited information available. It goes beyond the standard rational expectations model of how traders should (in theory) behave. It will serve as our launch-point into the development of a multi-trader model for describing financial markets, hence we will discuss it in some detail.

Suppose there is a popular bar—which Arthur envisaged as the El Farol bar in Santa Fe (United States)—and that you have to decide whether to show up on a given night each week. This could be a Friday night, say, when your favourite band is playing. Your goal is to attend *provided* you can get a seat. The trouble is, there are about $N = 100$ other people all trying to decide the same thing—there is only limited seating (say for $L_{\text{bar}} = 60$ people) and we assume that all potential attendees wish to sit down. Hence, it will not be worth the effort of showing up if the bar is overcrowded. But you do not know whether it will be overcrowded without phoning the other 100 people and asking them what they plan to do. And there are two problems with such phoning around: you may not know how many other people are potential customers, or their names or telephone numbers, and even if they told you, should you believe them? Let us imagine that, each week, the bar manager publishes the actual number of attendees, N_{attend} in the local paper the following day. Hence, everyone knows the previous outcomes. Based on this, you each try to predict whether the bar is likely to be *overcrowded* the following week (implying you should stay at home, running the risk that the bar is actually undercrowded) or *undercrowded* (implying you should make the effort to show up, running the risk that the bar is actually overcrowded). The $N = 101$ potential customers (including you) quickly realize that predictions of how many will attend depend on others' predictions of how many will attend (since this determines the actual attendance). But others' predictions in turn depend on their predictions of others' predictions. And so on. Hence, there is no correct expectational model—if everybody made the same decision, it would automatically be the wrong decision since everyone would either stay away (in which case you should have turned up) or they all show up (in which case you should have stayed away). A so-called mean-field theory describing a 'typical' agent will not work.

[6] See Arthur, W. B. (1999) *Science* **284**, 107.

This situation can be simulated on a computer choosing a pool of possible prediction methods given a particular set of recent outcomes, and randomly assigning a few such prediction methods to each agent. In this way, we capture the idea that the market of customers is heterogeneous, and that the quality of a given strategy depends on the strategies being used by all other agents. The practical problem with such a simulation is that it is difficult to define a suitable pool of strategies. Since the actual attendance numbers range from 0 up to N, the number of possible patterns of past attendances is enormous and hence so is the number of possible strategies. We will therefore abstract the basic features of this model using a binary approach, as described in more general terms in Section 4.4.

4.4 From the bar to the market

The El Farol bar problem of Section 4.3 alluded to a system where agents compete against each other for a limited resource: seating in a crowded bar. Each agent used only global information concerning past attendance numbers to work out whether to attend or not. Since the crucial ideas of competition, limited resources, adaptivity, etc., are well represented in the El Farol bar problem, we will explore it as a simplistic analogy to the action of traders in a financial marketplace. However, in making this analogy, we have to ask ourselves the following questions:

- What is the global information in a financial market?
- How do financial market agents (i.e. traders) decide how to trade?
- How do financial market agents 'win'?
- What other important properties do real financial market agents possess, that is, what else is missing?

We will discuss each of these questions in turn with the aim of forming a model of the behaviour of financial market traders in the spirit of the El Farol bar problem.

4.4.1 What is the global information in a financial market?

There are clearly many types of information available to traders: for example, the price histories of assets, histories of traded volumes, dividend yields, market capitalization, recent news, company reports. In reality all, some, or none of these information sources may actually be useful in making an investment decision for a particular asset. However, it is not our interest here to work out which of these information sources is *actually* useful; we simply need to know which sources financial market agents tend to use most. This question is easier to answer, and underlies the attractiveness of modelling the traders rather than attacking the market dynamics head-on. If we think about what we *ourselves* see most of, in relation to financial market assets such as stock, the answer has to be its price. The media is full of reports on recent price behaviour. Similarly, charts of prices occupy the majority of traders' screens on trading floors. It thus seems reasonable to take the source of global information upon which the traders act, to be based on the past history of prices for the financial asset of interest.

We now have to decide on a method of 'encoding' this past history of price movements in a simple yet representative way. The simplest alphabet we can use to describe the past history of an asset's price, is the binary alphabet of 0s and 1s. The simplest encoding would then be to assign one letter of this alphabet to each timestep.[7] For example, in the El Farol bar problem, one could assign a '0' to the state where the bar was undercrowded (more seats than customers) and a '1' to the overcrowded state. In this way, an agent in the El Farol bar problem could look at the string of recent outcomes (e.g. . . . 101110) and work out which nights it would have been best to attend. In a similar vein, the past history of asset prices $x[t]$ can be encoded by assigning a '0' to a price movement $\Delta x[t,\ t-1] = x[t] - x[t-1]$, which is smaller than a given value $L[t]$ (i.e. $\Delta x[t,\ t-1] < L[t]$) and a '1' for $\Delta x[t,\ t-1] > L[t]$. In this way the global information represents a caricature of the past history of asset prices with relation to the quantity $L[t]$. In the El Farol bar problem, the resource level $L[t]$ is related to the seating capacity of the bar against which attendance is judged. In the case of a financial market, $L[t]$ could represent a number of financial or economic variables, which could be either endogenous to the market or exogenous. An endogenous example could be the movement of a second financial product against which the movement of the asset of interest is judged, for example, the market index. An exogenous example could be the arrival of external news. Changing $L[t]$ affects the system's quasi-equilibrium, hence $L[t]$ could also be used to mimic a changing external environment due to some macroeconomic effect: for example, if interest rates are low, people may be tempted to put their money into the stock market. Conversely, if interest rates become high, then people may seek the risk-free choice of a high-interest bank account. In short, $L[t]$ indicates some measure of the attractiveness of the stock, or stock market as a whole, just as it indicates the attractiveness of the bar in the El Farol bar problem.

Since we wish to consider agents with limited capabilities, or equivalently agents who only deem recent information relevant, we shall build our source of global information from only the m most recent outcomes (0s and 1s). This has the immediate

[7] The idea of encoding global outcomes in multi-agent games as binary digits, and the subsequent discussion of the resulting strategy space, is due to Challet, D. and Zhang, Y. C. (1998) *Physica A* **256**, 514. See www.unifr.ch/econophysics/minority for a complete bibliography.

consequence that the number of possible states for the global information is finite and equal to $P = 2^m$. For $m = 2$, the only possible patterns in price are: up–up, up–down, down–up, and down–down (where 'up' and 'down' are with reference to $L[t]$). We can thus think of the different possible states of the global information variable forming a 'space' with one state for each unique asset price history (e.g. 11, 10, 01, and 00 for $m = 2$). We refer to this space as the 'history space' and denote the state at time t as the decimal equivalent of the string of m 0s and 1s: $\mu[t] \in \{0 \ldots P - 1\}$. For example, the history bit-string 00 corresponds to $\mu = 0$ while 11 corresponds to $\mu = 2^2 - 1 = 3$. The dynamics of the model within this history space can be represented on a directed graph. The particular form of directed graph relevant to this model is called a de Bruijn graph. Examples are shown in Fig. 4.2.

We have addressed what represents global information, but we have not discussed how it is generated. This amounts to the question 'what makes an asset price rise or fall in a financial market?' It is commonly believed that excess demand—that is,

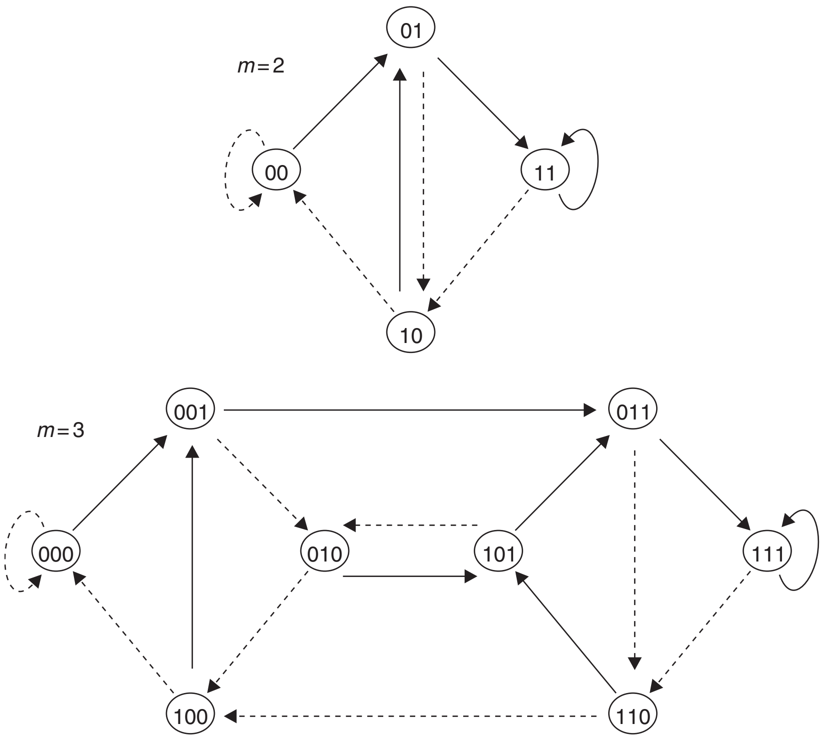

Fig. 4.2 History space for $m = 2$ and 3, showing the states available and transitions possible in the global information variable $\mu[t]$. Dotted line transitions imply negative price increments with respect to $L[t]$ (i.e. most recent outcome is '0'), whereas solid line transitions imply positive price increments with respect to $L[t]$ (i.e. most recent outcome is '1').

the difference between the number of assets sought and the number offered—exerts a force on the price of the asset. Furthermore, it is believed that a positive excess demand will force the price up and a negative demand will force the price down. A reasonable first-order approximation to the price formation process would then be:

$$\ln(x[t]) - \ln(x[t-1]) = D[t^-]/\lambda \tag{4.1}$$

or

$$x[t] - x[t-1] = D[t^-]/\lambda, \tag{4.2}$$

where $D[t^-]$ represents the excess demand in the market just *prior* to time t, while time t represents the time when the new price $x[t]$ is set and the buy/sell orders are executed.[8] The scale parameter λ represents the 'market depth', that is, how sensitive a market is to an order imbalance. In general, we would expect λ to be some increasing function of the number of traders N trading in that asset. If there are only a few agents trading the asset, the impact of each of them on the price is likely to be higher than if there are many agents, since the market-maker will then attempt to make the same amount of money off fewer agents. Numerical studies of several markets have shown a linear relationship between order imbalance and price-change to be broadly appropriate.

To summarize, we will take the global information upon which agents (traders) base their investment decisions to be the past history of the m most recent price movements $\Delta x[t, t-1]$ relative to $L[t]$. Each price movement $\Delta x[t, t-1] = x[t] - x[t-1]$ is taken to be a function of the excess demand in the market just prior to t (i.e. the demand which has built up between $t-1$ and t). This excess demand is determined by the aggregate of the agents' individual investment decisions. We thus have a feedback process: global information dictates investment decisions, investment decisions create a demand, total demand then generates new global information, and so on. This strong feedback is an essential feature of financial markets and as such should be considered as one of the essential ingredients for a market model.

4.4.2 How do financial market agents decide how to trade?

Our financial market agents act inductively[9] in order to make their investment decisions, just as in the El Farol bar problem. The only information an agent has available is the global information variable $\mu[t]$ and any parameters internal to the agent himself, such as his own wealth or investment track record. Real financial market

[8] Typically the price-changes at each timestep will be small compared to the price. Hence, $\ln(x[t]) - \ln(x[t-1])$ is approximately proportional to $x[t] - x[t-1]$ implying Equations (4.1) and (4.2) have the same basic form.

[9] There is no a priori 'correct' action that they can deduce.

agents operate on similar grounds. A reasonable representative model of a financial market agent might be someone with preconceived ideas of which way to trade given a particular set of circumstances. The agent may modify which of these preconceived ideas he trusts most, based on its previous successes or failures at predicting which investment decision should be made. Let us therefore imagine that each agent owns a small library of s strategy books. A strategy book details what investment decision the agent should make, based on the observed recent price pattern which is represented by the global information $\mu[t]$. As such, the book details chartist principles. We denote the suggested action (investment decision) of strategy book R, given global information (previous price pattern) $\mu[t]$, to be $a_R^{\mu[t]}$. If a strategy book is 'full', then it will have a suggested action for each possible state of the global information. Thus, each of these strategy books has 2^m pages, one for each pattern. If there are only two possible suggested actions (e.g. 'buy' and 'sell') there will be 2^{2^m} possible books an agent can buy, and subsequently use as guidelines on how to trade. In other words, an agent in possession of one of these strategy books, whichever he chooses to buy, can find a page giving guidelines on how to trade in every possible market state. It seems unlikely however, that in reality such a complete book exists for any general value of m. Indeed, even if such a book did exist, it is unlikely that a given market participant would consider all the pages to be useful. For example, page one of a strategy book may say that if the asset price has fallen three days running one should sell on the fourth day. Page three may say that if the asset price has fallen, then recovered and then fallen again, one should buy. The agent holding this book may well believe in page one but think that the guidelines of page three are useless; he himself considers the pattern down–up–down to be no trading signal at all. This agent would therefore continue to hold any previous position he had if he saw the pattern down–up–down. Nevertheless, the *set* of strategy books is complete. By this we mean that every possible (binary) response to every possible (binary) price history is represented. Agents may seek to trade in completely different ways to each other, but as long as the behaviour of the agents is characterized by taking a binary action based on binary global information, then every possible set of beliefs and preconceptions that an agent could have can be ascribed to holding[10] a certain (limited) library of these 2^{2^m} strategy books. Of course, agents will in general have quite different libraries—equivalently, they will

[10] An agent does not necessarily have to 'buy' such strategy books. These books could also represent a set of innate rules that an agent may already carry in his head, resulting from his individual character and temperament or background, or resulting from the trading restrictions placed on him by the financial institution for which he works. Hence, instead of 'buying' certain strategies, the agent can be thought of as 'buying into' certain strategies. The effect is the same. As a by-product of this approach, we note that we are indirectly incorporating the basic notion of 'behavioural finance'.

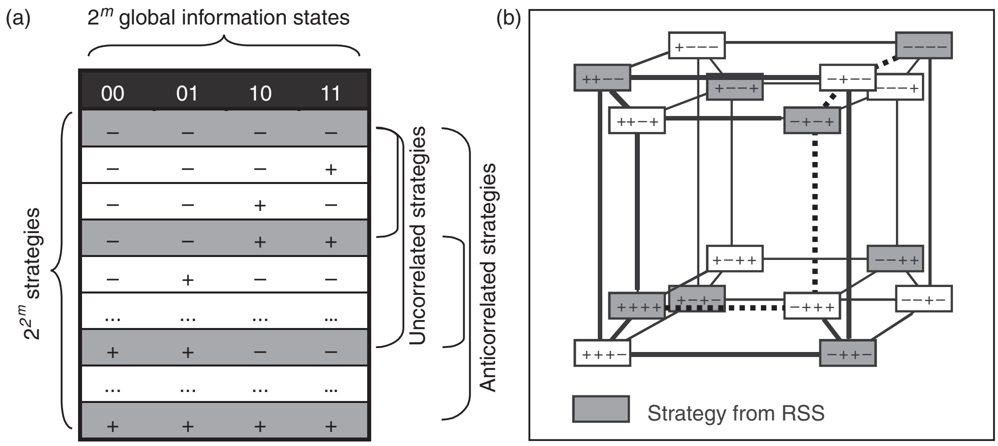

Fig. 4.3 'Strategy space' for $m = 2$. (a) Schematic representation of the $2^{2^m} = 16$ different strategies (i.e. strategy books). The greyed strategies belong to the Reduced Strategy Space (RSS) and are either totally uncorrelated or anticorrelated with respect to each other. There are $2^{m+1} = 8$ strategies in the RSS. (b) Representation of a $2^m = 4$ dimensional hypercube, which demonstrates the Hamming distance between strategies. The minimum number of edges linking strategies is the Hamming distance; for example, the dotted line shows a Hamming distance of 4 between strategies $----$ and $++++$.

have bought different books[11]—hence they will *not* take the same actions in the same circumstances. In other words, agents are heterogeneous—just what is expected from real life, but just what is *not* included in standard economic models.

The beauty of dealing with a complete set of binary strategy books (i.e. strategies) is that the 'space' of strategies can be broken down in a logical and useful manner. In Fig. 4.3, we show a picture of the strategy space for the case $m = 2$, using $\{-,+\}$ to denote the two possible actions $\{-1,+1\}$ for each global information state $\mu[t]$.

Strategies within the subset which is shaded grey, possess particularly simple interrelationships.[12] In particular, pairs of strategies taken from this subset have one of the following types of correlation:

Anticorrelated. For example, any two agents using the strategies $----$ and $++++$, respectively, would take the opposite action irrespective of the sequence of previous outcomes, that is, irrespective of the global information state. Hence, one agent will *always* do the opposite of the other agent. For example, if one agent buys at a given timestep, the other agent will sell. Their net effect on the excess demand $D[t]$ therefore cancels out at each timestep. Hence, together they will not contribute to fluctuations in $D[t]$ and hence the price. In short, *they do not contribute to the volatility*.

[11] See note 10. [12] See note 7.

This is a crucial observation for understanding the behaviour of the volatility in this system, as we discuss later.
Uncorrelated. For example, any two agents using the strategies $----$ and $--++$, respectively, would take the opposite action for two of the four histories, while they would take the same action for the remaining two histories. Assuming that the $m = 2$ histories occur equally often, the actions of the two agents will be uncorrelated on average.

Although there are $2^P \equiv 2^{2^{m=2}} = 16$ strategies (i.e. strategy books) in the strategy space, there are subsets of strategies that can, therefore, be classed as either purely anticorrelated or uncorrelated.[13] Consider, for example, the two groups:

$$U_{m=2} \equiv \{----,\ ++--,\ +-+-,\ -++-\},$$
$$\overline{U_{m=2}} \equiv \{++++,\ --++,\ -+-+,\ +--+\}.$$

Any two strategies within $U_{m=2}$ are uncorrelated since they have a relative Hamming distance[14] of $P/2$. Likewise, any two strategies within $\overline{U_{m=2}}$ are uncorrelated since they have a relative Hamming distance of $P/2$. However, each strategy in $U_{m=2}$ has an anticorrelated strategy in $\overline{U_{m=2}}$: for example, $----$ is anticorrelated to $++++$, $++--$ is anticorrelated to $--++$, etc. This subset of strategies comprising $U_{m=2}$ and $\overline{U_{m=2}}$, forms a Reduced Strategy Space (RSS).[15] Since it contains the essential correlations of the Full Strategy Space (FSS), it turns out that running a simulation within the RSS reproduces the main features of the full model. The RSS has a reduced number of strategies, $2.2^m = 2P \equiv 2^{m+1}$, as compared to the FSS which has $2^P \equiv 2^{2^m}$. For $m = 2$, there are 8 strategies in the RSS as compared to 16 in the FSS. The choice of RSS is not unique, that is, a given FSS can have many possible choices for RSS.

4.4.3 How do financial market agents 'win'?

Each agent holds a library of s strategy books, which map the present available global information $\mu[t]$ to an action $a_R^{\mu[t]} \in \{-1, +1\}$. Let -1 indicate a sell decision and $+1$ indicate a buy decision. It is essential that the agents should hold $s > 1$ strategy books in their library in order that they may adapt their behaviour to the current market conditions. An agent with only one strategy book available would have no other choice

[13] There is the additional, trivial case of 'fully correlated' which represents the correlation of a given strategy with itself.

[14] A convenient measure of the distance (i.e. closeness) between any two strategies is the 'Hamming distance', defined as the number of bits that need to be changed in going from one strategy to another. For example, the Hamming distance between $----$ and $++++$ is $P = 4$, while the Hamming distance between $----$ and $--++$ is just $P/2 = 2$.

[15] See note 7.

than to follow blindly the only investment suggestion available, that is, follow his one strategy book. An agent with $s = 1$, therefore, does not have the opportunity to learn how to trade most effectively. The crucial question is then 'How do the agents choose which of their s strategy books to use in order to trade effectively?'. One physically realistic answer is as follows: each agent uses the strategy book (i.e. strategy) which would have been the most successful in the past, judging from the past history of the market. To work this out, each agent needs to keep a tally of the success rate $S_R[t]$ for each of his strategies (R is the strategy label). Hence, different agents holding the same strategy book will agree on its relative merit. This is a key feature of the model design: it can lead to large groups of agents agreeing which is the best strategy book in their separate libraries and thus *independently* making the same investment decision. The consequence of this is that agents might rush to the market in groups (crowds) even in the absence of any direct communication between each other.

We still need to define exactly how the agents will judge the success of, and hence reward, their strategy books. In the El Farol bar problem, there are N regular bar-goers but only $L_{\text{bar}} < N$ seats. The reward scheme is therefore simple: bar-goers are successful if they attend *and* they manage to obtain a seat. Thus, any strategy book that suggests the agent go to the bar at time $t(a_R^{\mu[t]} = 1)$ when it turns out the bar is undercrowded ($N_{\text{attend}}[t] < L_{\text{bar}}$), should be rewarded positively. A suitable strategy reward structure could thus be proposed as:

$$S_R[t+1] = S_R[t] - a_R^{\mu[t]}(N_{\text{attend}}[t] - L_{\text{bar}}), \tag{4.3}$$

where $N_{\text{attend}}[t]$ is the number of attendees at time t. This reward system gives the strategy books a pay-off proportional to how under/overcrowded the bar actually was. The reward structure of Equation (4.3) can be generalized to:

$$S_R[t+1] = (1 - 1/T)S_R[t] - a_R^{\mu[t]}\chi[N_{\text{attend}}[t] - L_{\text{bar}}], \tag{4.4}$$

where T represents a timescale over which the previous successes or failures of the predictions of strategy book R are 'forgotten'. The function $\chi[x]$ in the second term is an odd, increasing function of x, which is typically taken to be either $\chi[x] = x$ or $\chi[x] = \text{sgn}[x]$. We emphasize that this reward scheme implicitly assumes that the bar-goers value sitting down above other criteria. While this makes sense for many situations, it is not universal. For example, customers of college-based bars do not seem to view seating as a necessary requirement for a successful evening. The motto 'the more the merrier' often seems more appropriate. In more general and realistic situations, the correct reward scheme is likely to be less simple than in Arthur's El Farol bar problem. Until a specific reward scheme is defined, the bar model remains ill-specified. Putting this another way, the precise reward scheme chosen is a fundamental property of the resulting system and directly determines the resulting dynamical evolution.

As we have discussed, the El Farol bar problem can be seen as being somewhat analogous to a financial market where the bar-goers are replaced by traders. In the same way that a general bar problem requires a non-trivial reward scheme, any realistic agent-based market model will need a non-trivial reward scheme in order to avoid inconsistencies with financial market microstructure. With this in mind, we will now pursue a more concrete connection between the bar problem and a market model, in terms of the strategy reward structure. Earlier we suggested an analogy between the attendance of bar-goers to El Farol, $N_{\text{attend}}[t]$, and the financial asset price change. These quantities provide the global information variables in each case. Taking Equation (4.4) and using this analogy, we can write down a strategy book pay-off $g_R[t+1]$ such that $S_R[t+1] = (1 - 1/T)S_R[t] + g_R[t+1]$. This implies $g_R[t+1] = -a_R^{\mu[t]}\chi[\Delta x[t+1,\ t] - L[t+1]]$. Using our price formation process from Equation (4.2), we then obtain:[16]

$$g_R[t+1] = -a_R^{\mu[t]}\chi[D[(t+1)^-]/\lambda - L[t+1]]. \tag{4.5}$$

The actions $a_R^{\mu[t]}$, when followed by the agents, represent buy or sell orders in the market. Hence the total demand is given by:

$$D[(t+1)^-] = \sum_{R=1}^{2^P} n_R[t]a_R^{\mu[t]}, \tag{4.6}$$

where $n_R[t]$ is the number of agents who are using strategy book R at time t in order to make their investment decisions. Strategy R is the most successful strategy book that these $n_R[t]$ agents hold.

We can see that timing is becoming important, with some quantities evaluated at $t+1$ and others at t, etc. So let us look more closely at these timings. Equation (4.6) is for the excess demand at time $(t+1)^-$, that is, just prior to time $t+1$. Thus, $D[(t+1)^-]$ can only result from all the information that is available at time t, that is, the global information $\mu[t]$ and the set of strategy book scores $\{S_R[t]\}$. From this information, the agents take actions $a^{\mu[t]}$ producing a total demand $D[(t+1)^-]$. The agents' actions (i.e. orders) however do not get realized (i.e. executed) until time $t+1$, when the new price $x[t+1]$ is known. The pay-off function Equation (4.5) rewards agents positively for deciding to sell (buy) assets, $a^{\mu[t]} = -1$ (1), when the number of buyers (sellers) minus sellers (buyers) exceeds $\lambda L[t+1]$. If we take the reference

16 The seating capacity L_{bar} in the bar problem, and the resource level $L[t]$ in the market model, are not numerically equivalent. This is because $\langle D[t]\rangle \approx 0$ for an unbiased market, while $\langle N_{\text{attend}}[t]\rangle \approx L_{\text{bar}}$ for the bar.

asset (e.g. a bond) to have a zero return, that is, $L[t+1] = 0$, we can see that the strategy reward structure of Equation (4.5) rewards strategy books for suggesting the *minority* trading decision, that is, to sell (buy) when there is a majority of buyers (sellers). The binary model with this assumption of $L[t] = 0$ for all t, is thus referred to as the 'Minority Game', which was originally proposed by Challet and Zhang. The structure of the Minority Game was further generalized by Johnson *et al.* to incorporate a variable number of active agents.[17] The agents in this generalized model, known as the 'Grand Canonical Minority Game' (GCMG),[18] each has a 'confidence level' r and will only participate in the game if the score of their best strategy book S_{R^*} is higher than this level (i.e. $S_{R^*} > r$). This feature whereby traders only participate (i.e. trade) when they are sufficiently confident of success, is a crucial ingredient for building a successful multi-agent market model.

Let us investigate why the goal of trading in the minority group could be a physically reasonable ambition for the agents. We introduce a notional wealth W_i of an agent i as follows:

$$W_i[t] = \phi_i[t]x[t] + C_i[t], \tag{4.7}$$

where ϕ_i is the number of assets held and C_i is the amount of cash held. It is clear from Equation (4.7) that an exchange of cash for assets at any price does not in any way affect the agents' notional wealth. However, the point is in the terminology: the wealth $W_i[t]$ is only *notional* and not real in any sense. The only *real* measure of wealth is C_i, the amount of capital the agent has available to spend. Thus, it is evident that an agent has to do a 'round trip' (i.e. buy (sell) an asset then sell (buy) it back) to discover whether a *real* profit has been made. Let us consider two examples of such a round trip: in the first case, the agent trades with the minority decision, and in the second he trades with the majority decision.

- trading with the minority decision

t	Action $a[t]$	$C_i[t]$	$\phi_i[t]$	$x[t]$	$W_i[t]$
1	submit buy order	100	0	10	100
2	buy..., submit sell order	91	1	9	100
3	sell	101	0	10	101

[17] Johnson, N. F., Hart, M., Hui P. M., and Zheng, D., (2000) *Int. J. Theo. Appl. Fin.* **3,** 443; Jefferies, P., Johnson, N. F., Hart, M., and Hui, P. M., (2001) *Eur. Phys. J. B* **20**, 493.

[18] See note 17.

- trading with the majority decision

t	Action $a[t]$	$C_i[t]$	$\phi_i[t]$	$x[t]$	$W_i[t]$
1	submit buy order	100	0	10	100
2	buy. . ., submit sell order	89	1	11	100
3	sell	99	0	10	99

As can be seen, trading with the minority decision creates wealth for the agent on performing the necessary round trip, whereas trading with the majority decision loses wealth. However, if the agent had held the asset for a length of time between buying it and selling it back, his wealth would also depend on the rise and fall of the asset price over the holding period. So, although the Minority Game strategy-book reward mechanism seems perfectly reasonable for a collection of traders who simply buy/sell on one timestep and sell/buy back on the next, this is *not* of course what real financial traders do in general. This is the main criticism of the Minority Game as a market model.

To keep consistency with the real financial market, it seems that we need a more subtle reward mechanism than simply rewarding strategy books which suggest trading with the minority decision. We have identified that whilst trading with the minority decision can be beneficial, the minority pay-off structure makes no consideration of the rise or fall in the value of the agent's portfolio ϕ_i of assets. Let us try to rectify this by examining the form of the agent's notional wealth, Equation (4.7). If we differentiate the notional wealth, we get an expression for $\Delta W_i[t+1,\ t] = W_i[t+1] - W_i[t]$:

$$\Delta W_i[t+1,\ t] = \Delta C_i[t+1,\ t] + x[t+1]\Delta\phi_i[t+1,\ t] + \phi_i[t]\Delta x[t+1,\ t].$$

The first two terms cancel because the amount of cash lost $-\Delta C_i[t+1,\ t]$ is used to buy the extra $\Delta\phi_i[t+1,\ t]$ assets at price $x[t+1]$. This leaves us with:

$$\Delta W_i[t+1,\ t] = \phi_i[t]\Delta x[t+1,\ t]. \tag{4.8}$$

We can then use Equation (4.8) to work out an appropriate reward $g_R[t+1]$ for each strategy based on whether its investment suggestion $a_R^{\mu[t]}$ would have induced a positive or negative increase in notional wealth. Let us first use the fact that the price change $\Delta x[t+1,\ t]$ is roughly proportional to the excess demand $D[(t+1)^-]$: this can be seen explicitly from our earlier equation for the price formation, Equation (4.2). We therefore obtain from Equation (4.8):

$$\Delta W_R[t+1,\ t] \propto \phi_R[t]D[(t+1)^-].$$

We then identify the accumulated position in the asset for strategy R at time t to be $\phi_R[t]$: this represents the sum of all the actions (investment suggestions) made by that strategy, which would have been executed between time 0 and t had the strategy been used. Remembering that at time t the action (order) $a_R^{\mu[t]}$ has not yet been executed (it gets executed at $t+1$), this gives $\phi_R[t] = \sum_{i=0}^{t-1} a_R^{\mu[i]}$. Let us then set the pay-off $g_R[t+1]$ given to a strategy book R, to be an increasing (odd) function χ of the notional wealth increase $\Delta W_R[t+1, t]$ for that strategy. We thus arrive at:

$$g_R[t+1] = \chi\left[\sum_{i=0}^{t-1} a_R^{\mu[i]} D\left[(t+1)^-\right]\right].$$

Note that if agents were comparing this notional wealth increase with the wealth increase available through investment in another asset (e.g. a bond) with return $L[t+1] \neq 0$, then we would have instead

$$g_R[t+1] = \chi\left[\sum_{i=0}^{t-1} a_R^{\mu[i]} \left[D\left[(t+1)^-\right]/\lambda - L[t+1]\right]\right]. \tag{4.9}$$

We could also propose a locally weighted equivalent of Equation (4.9) where the reward given to a strategy is more heavily weighted on the result of its recent actions, rather than the actions it made further in the past. This gives:

$$g_R[t+1] = \chi\left[\sum_{i=0}^{t-1} (1-1/T)^{t-1-i} a_R^{\mu[i]} \left[D\left[(t+1)^-\right]/\lambda - L[t+1]\right]\right], \tag{4.10}$$

where T again represents a characteristic timescale over which the position accumulated by the strategy is 'forgotten'. In the limit $T = 1$, Equation (4.10) becomes $g_R[t+1] = \chi[a_R^{\mu[t-1]}\left[D\left[(t+1)^-\right]/\lambda - L[t+1]\right]$: that is, only the position resulting from the most recently executed trade is taken into account. With $T = 1$, this pay-off structure essentially rewards a strategy at time $t+1$ based on whether the notional wealth change $\Delta W_R[t+1, t]$ was more positive than it would have been if action $a_R^{\mu[t-1]}$ had not been taken.

If Equation (4.10) is used in an agent-based market model, the agents play the strategy they hold which has collected the highest 'virtual' notional wealth. We mean 'virtual' in the sense that the strategy itself will not have actually collected this notional wealth, unless it has been played incessantly since time $t = 0$. The agents in this model are hence all striving to increase their notional wealth, and are allowed to do so by taking arbitrarily large positions ϕ_i. This suggests that the agents of a realistic market model need more attributes than simply a library of strategy books in order to function in a realistic manner: after all, real financial agents have finite resources and cannot take such arbitrarily large positions.

4.4.4 What else is missing?

In the El Farol bar problem, the agents have to make the simple decision as to whether or not attend the bar. In reality, such decisions may also be influenced by an agent's available funds, or maybe even whether they are still hung-over from the last time they visited. While such complications could be incorporated, this would deviate from our quest for a minimal model, that is, a model that explains the observed market behaviour yet employs a minimal number of physical characteristics. However, in order to generate stable and realistic market dynamics, we will see that *some* additional properties of 'real' traders do indeed need to be incorporated. Let us start by looking at the dynamical behaviour of the GCMG: this is the case of the El Farol bar problem with $L_{\text{bar}} = N/2$ (which corresponds to $L[t] = 0$) but generalized such that agents only participate if the score of their perceived best strategy book is greater than the confidence level r. We will compare the GCMG's dynamical behaviour with that of the corresponding financial market model (MM) constructed with a strategy pay-off structure given by Equation (4.10) with $L[t] = 0$. We have simulated each with the same parameters: $N = 501$ agents, $s = 2$ strategies per agent, $T = 100$ time horizon beyond which strategy scores 'forget' past performance, an $r = 4$ pay-off point confidence-to-trade threshold, and the $\chi = \text{sgn}$ strategy pay-off ($+1/-1$ point for a good/bad investment suggestion). Each simulation was run for agents with low memory $m = 3$ and for agents with high memory $m = 10$. The resultant price $x[t]$ and number of active agents $V[t] = \sum_{R=1}^{2^P} n_R[t-1]$ are displayed in Fig. 4.4. Note that both the price-change and volume at time t depend on the agents' actions at time $t - 1$.

The two different values of the agent's global information memory length m (or more formally the ratio P/N) in Fig. 4.4, describe markets in two distinctly different 'phases'. Arguably one of the most important features of these types of model is that there is a common perception of a given strategy R's (virtual) success. For example, $S_{R=1}[t]$ for one agent is the same as $S_{R=1}[t]$ for another agent. This implies that the number of agents adopting the same action is dependent on the number which hold each strategy. At low memory (low P/N) there are relatively few strategies available but many agents; consequently many agents will hold the same strategy. This implies that in this 'crowded phase', there will be a large group of agents using a successful strategy R^* and hence adopting the same action $a_{R^*}^{\mu[t]}$ (i.e. a *crowd*). However, there will only be a small group using a strategy such as the anticorrelated strategy $\overline{R^*}$ and hence adopting the opposite action $a_{\overline{R^*}}^{\mu[t]} = -a_{R^*}^{\mu[t]}$ (i.e. an *anticrowd*). The result of this high agent coordination is a volatile market with large asset price movements. Conversely for high m (high P/N), hardly any agents hold the same strategy and so the crowds and anticrowds are small and of similar size. This results in a market with low agent coordination and consequently lower volatility and fewer large changes in

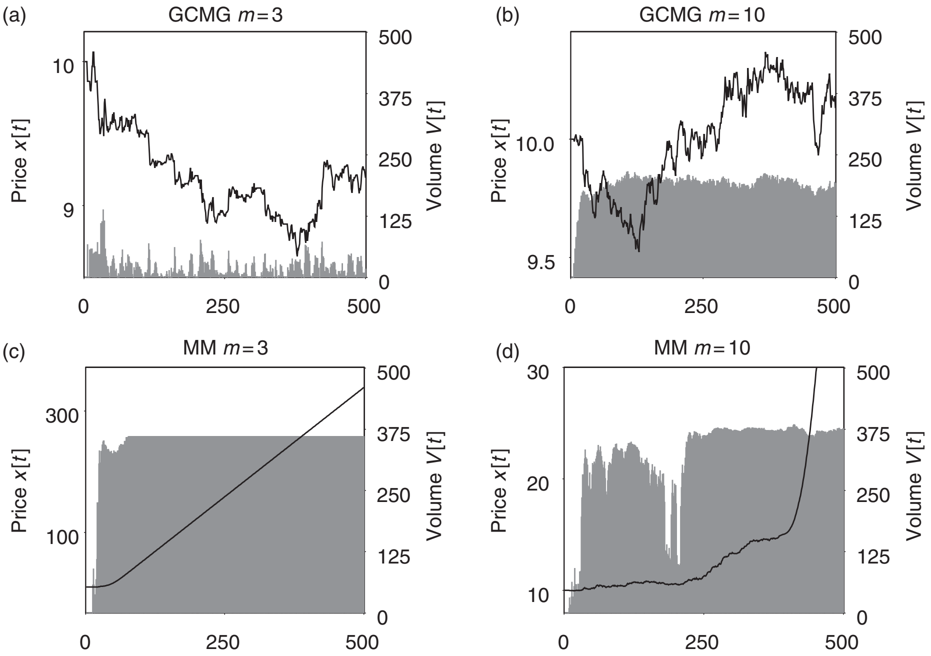

Fig. 4.4 Temporal evolution of the price $x[t]$ and 'volume' (i.e. number) of active agents $V[t]$ in the (a and b) GCMG and the (c and d) Market Model (MM) for two different values of the global information memory length m. Parameters for the simulations were: $N = 501, s = 2, T = 100$, and $r = 4$ with binary (i.e. $\chi = \text{sgn}$) strategy pay-off.

the price. With this distinction between low and high m regimes in mind, let us briefly describe the dynamics in each model.

GCMG. At low m and with a sufficiently high confidence-to-trade threshold, the GCMG is able to reproduce many of the stylized facts of real markets. It can be seen from Fig. 4.4 that the number of active agents (volume) is generally low and bursty. The asset price series is thus characterized by frequent large movements, giving fat-tailed distributions of returns, and clustered volatility. The autocorrelation of the GCMG asset price movements is essentially zero: as soon as lots of agents start taking the same action, the strategies which produce that action are penalized as can be seen from Equation (4.5). At high m, the absence of agent coordination leads to a lack of activity clustering, hence the series of asset prices appears more random.

MM. At low m, the MM asset price very quickly acquires a steady trend. Any agent with at least one trend-following strategy (i.e. price changes $\ldots \Delta x[t-2] > 0$, $\Delta x[t-1] > 0$, $\Delta x[t] > 0 \Rightarrow a_R = 1$) will join the trend and hence benefit, notionally, from the consequent asset price movement. Because the strategies and hence agents are allowed to accrue limitless positions, the trend is self-reinforcing since ΔW_R just keeps getting bigger for the trend-following strategies. At high m, the lack of agent coordination means that it is harder for the model to find this attractor—however, sooner or later a majority of trend-following strategies will have collected sufficiently large scores S_R and positions $|\phi_R|$ to be traded successfully. From then on, the pattern of success is self-reinforcing and again the steady trend is created. This result is the natural consequence of wanting the agents to maximize notional wealth at the same time as being allowed arbitrarily large positions.

So it seems that whilst the dynamics of the adapted El Farol bar problem, the GCMG, are stable and market-like, the dynamics of the so-called MM are unstable and un-market-like. By trying to develop a strategy reward structure which more realistically models the actions of a financial market agent, as opposed to a bar attendee, we have introduced an instability into the model. This shows all too clearly how the strategy book reward structure defines the nature of the model itself. This instability in the MM seems to arise from the potential for the agents to take up arbitrarily large positions in the financial asset. This is clearly not a physically realistic situation: after all, nobody is infinitely rich, and only a finite amount of the asset is available. Even if the MM agents cannot *actually* buy (sell) more assets because the market has become illiquid, as long as there is still a positive (negative) *demand* for the asset then the MM price will keep rising (falling). Therefore, even by pretending to want to buy (sell) more assets, the agents can manipulate the price up (down) to profit their long (short) position. Therefore, the MM should have a mechanism which stops the agents from demanding to buy (sell) more assets. Perhaps the most obvious contender for this mechanism is the finite resources of the agents: an agent cannot place an order to buy (sell) assets if he does not have the required funds (assets) to complete the transaction. Let us then investigate whether adding the property of finite resources to the agents of our model, helps stabilize the dynamical behaviour.

4.4.4.1 Agent wealth. Financial market agents only have finite resources. The amount of resources available to the agents will depend on how wealthy they are, or how tight the regulations imposed by their risk managers are. For example, some institutions will not allow themselves to be over-exposed to the risky movements of a particular asset and so will insist that positions in this asset are limited to a certain number of quanta. This limitation has the effect of putting a ceiling on the demand to buy or sell assets in the marketplace: the hard limit of the agents' resources thus in turn imposes a hard limit on the magnitude of price trends. We can include the effect of limited agent resources in our market model (MM) by allocating agent i an initial capital $C_i[0]$, and position $\phi_i[0]$, and then updating this capital using:

$$C_i[t] = C_i[t-1] - \Delta\phi_i[t, t-1]x[t]. \qquad (4.11)$$

We can then further impose the limitations to trading based on the agents' inventory of cash and assets. Specifically, an agent is not permitted to trade at time t if:

$$a_R^{\mu[t]} = 1 \ \cap C_i[t] < x[t],$$
$$a_R^{\mu[t]} = -1 \ \cap \phi_i[t] < 0.$$

In words, the first condition states that an agent cannot submit a buy order unless he has at least enough capital to buy the asset at the quoted price. The second condition states that he cannot submit an order to short-sell if he already holds a short position. If we imposed no limit on short-selling, an unstable state of the system would exist wherein all agents short-sell indefinitely. We initialize this generalized MM market model (we will call it MM(W) where W represents 'wealth') with agent resources such that the agents' initial buying power is equal to their initial selling power, that is, $\{C_i[0], \phi_i[0]\} = \{n\, x[0], n-1\}$. Hence, each agent starts off with the power to buy or sell n assets (he is allowed up to one short-sell). If we then increase n from $n = 1$, we see a qualitative change in the dynamical behaviour of this market model with the periods of trending grow longer as n increases. We investigate this behaviour by fixing n and running the model for 1500 timesteps, recording the value of the global information $\mu[t]$ in the last 500 timesteps t. We then count the number of times within these last 500 timesteps when either $\mu = 0$ occurs (implying negative price movements over the last m timesteps) or $\mu = P - 1$ occurs (implying positive price movements). We denote the frequency with which these states occur as f_{trend}. If the model visited all states of the global information $\mu[t]$ equally, we would expect $f_{\text{trend}} = 1/P + 1/P = 2/P$: in particular, this would arise if the states were visited randomly.

Figure 4.5 shows the variation of f_{trend} with n, and has several interesting features. First it can be seen that as n is increased, and the consequent wealth available to the agents grows, the tendency of the model to be dominated by price trending increases dramatically. Also, we see that at low n, f_{trend} is below the limit of equally visited (e.g. randomly visited) μ states. This is due to the high degree of anti-persistence in the system, which arises because the demand is so tightly bounded by the agents' limited resources. The large spread in the results arises from the tendency of the model to exhibit clustering of states of activity: persistence follows persistence and anti-persistence follows anti-persistence. This activity clustering in turn arises because an agent's strategy book will be penalized for suggesting that an agent join a trend, since this corresponds to trading with the majority group. However, once the trend has been joined, the strategy book will be rewarded as the notional value of its position grows.

The major source of agent diversity in the market models so far, has been the heterogeneity in the allocation of strategy books in the agents' libraries. However, the

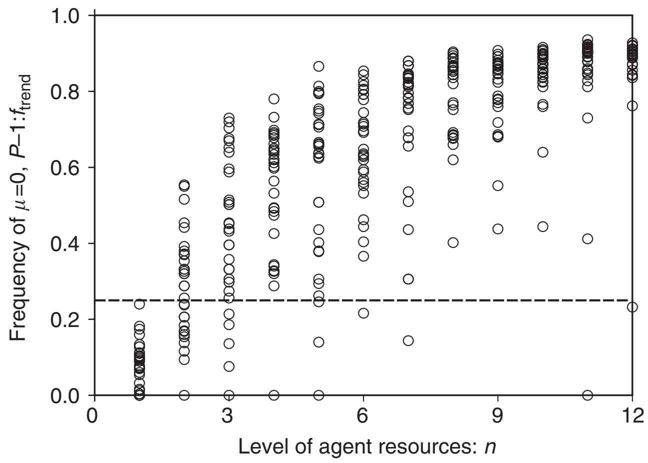

Fig. 4.5 The frequency of occurrence, f_{trend}, of the global information states $\mu = 0, P-1$ as a function of each agent's capital resource level n. The results were taken from the last 500 timesteps of a 1500 timestep simulation. The market model used was the MM(W) model with evolving agent wealth. The model parameters were $N = 501$, $m = 3$, $s = 2$, $T = 100$, and $r = 4$, with a binary ($\chi = \text{sgn}$) strategy reward scheme. The dashed line represents equally visited μ-states, hence $f_{\text{trend}} = 2/P = 2/2^3 = 0.25$.

introduction of an agent wealth brings about a secondary source of diversity. Even if we initiate the model with all agents having an equal allocation of wealth in the form of cash plus assets, the wealth of the agents W_i will soon become heterogeneous as a direct result of their heterogeneous strategies. Figure 4.6 shows the heterogeneity of agents' wealth growing with time during a MM(W) simulation. After many timesteps have elapsed, the distribution of agents' wealth seems to reach an equilibrium in which many agents have lost the majority of their wealth to a minority of agents—this minority now holds significant wealth. In other words, we are witnessing the spontaneous creation of a wealth-based hierarchy within the multi-agent society.

The heterogeneity of agents' wealth in the MM(W) model is fed back into the system through the buying power of the agents. However, although wealthier agents have the potential to buy and sell more assets, they still only trade in single quanta of the asset at any given timestep, just like the poorer agents. Therefore, it may be more reasonable to propose that the agents trade in sizes proportional to their individual wealth, that is,

$$\Delta\phi_i[t, t-1] = \gamma \frac{C_i[t-1]}{x[t-1]}, \quad \text{for} \quad a_{R_i^*}^{\mu[t-1]} = 1,$$
$$\Delta\phi_i[t, t-1] = \gamma \phi_i[t-1], \quad \text{for} \quad a_{R_i^*}^{\mu[t-1]} = -1,$$

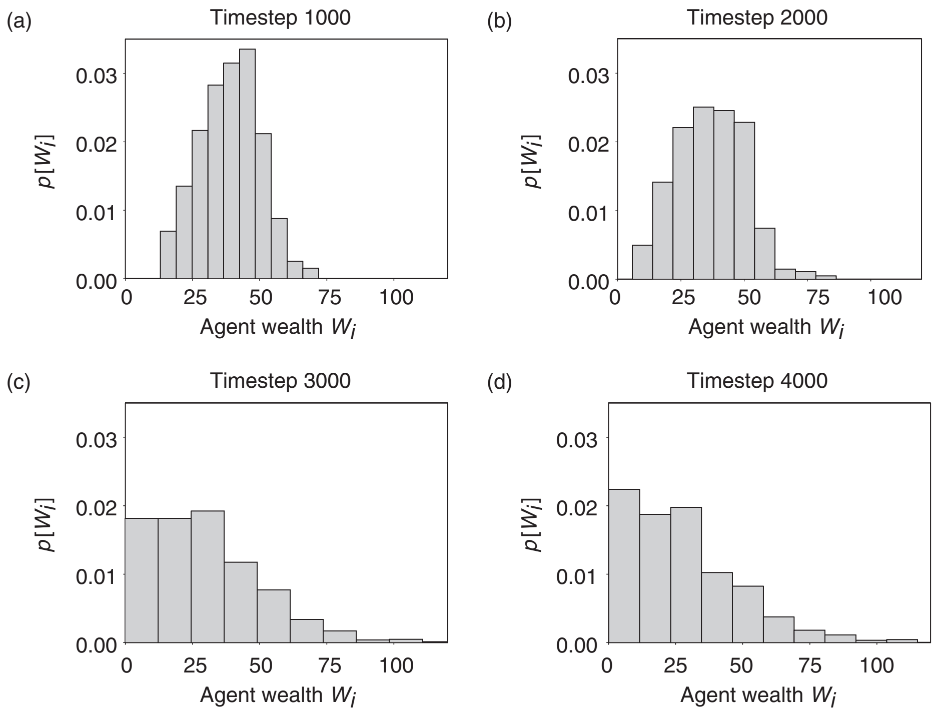

Fig. 4.6 Probability distribution of the agents' wealth W_i at four different times t during evolution of the MM(W) model. Initially all agents were allocated the same resources $\{C_i[0], \phi_i[0]\} = \{3x[0], 2\}$ with the initial price $x[0] = 10$. Parameters for the simulation were $N = 1001$, $m = 3$, $s = 2$, $T = 100$, and $r = 4$, with a binary ($\chi = \text{sgn}$) strategy reward scheme.

where R_i^* is the highest scoring strategy of agent i. The factor γ then enumerates what fraction of an agent's resources (i.e. cash for buying, assets for selling) he is willing to transact at any given time. In general, assets need to be divisible in this system, that is, the sense of agents buying and selling a quanta of the asset as in MM(W) is lost. This in turn means that instead of the degree of trending being controlled by the level of initial resource allocation n, it is instead determined by n/γ since this effectively determines the number of trades agents can make in any trending period before hitting the boundary of their capital resources. Also, with this system of trading in proportion to wealth, trends will start steep and end shallow as agents run out of resources and thus make smaller and smaller trades. Apart from these qualitative differences, the system has a very similar dynamical behaviour to the more straightforward MM(W) model.

Diversity in strategies and wealth are the two big sources of agent heterogeneity that we have covered so far. This agent heterogeneity has led to a market model

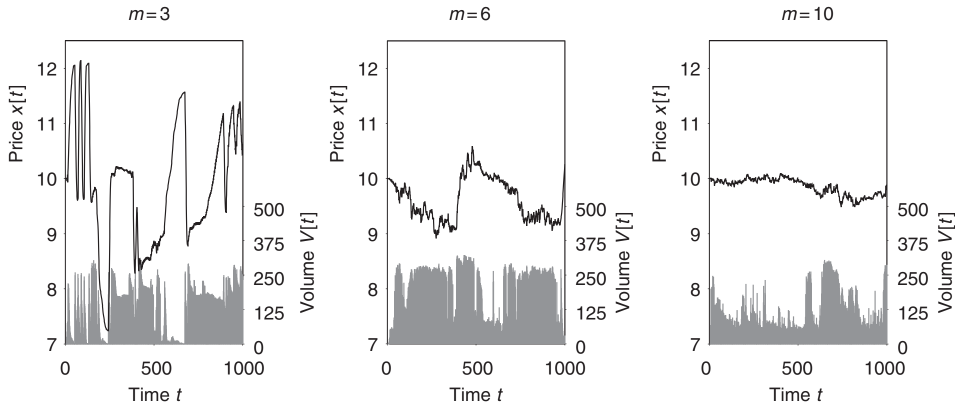

Fig. 4.7 Examples of the dynamical behaviour of the price $x[t]$ and volume $V[t] = \sum_{R=1}^{2^P} n_R[t-1]$ of the MM(W) model for three different levels of crowding, as determined by the memory length. The parameters for the simulation are $N = 501$, $m = 3$, $s = 2$, $T = 100$, and $r = 4$, with a binary ($\chi = \text{sgn}$) strategy reward scheme.

with dynamical behaviour that is interesting and diverse over a large parameter range. However, the typical price/volume output only starts to become representative of a real financial market at higher m, as can be seen in Fig. 4.7. As discussed earlier, the high m regime represents a 'dilute' market where very few traders act in a coordinated fashion. However, we would expect that a real financial market is *not* in a dilute phase at all, but instead has large groups of agents forming crowds which rush to the market together creating a bursty pattern of activity. Why then do these models, when pushed into the smaller m regime as shown in Fig. 4.7, produce endless bubbles of positive followed by negative speculation?

4.4.4.2 Trading timescales. To answer the above question, we investigate further the subject of agent diversity. Although our agents have differing sets of strategies and consequently different wealth, they all act on the same timescale. When we ourselves look at charts such as Fig. 4.7, we see patterns not only on a small point-to-point scale but also on a much larger scale. In short, we try to identify patterns over a wide range of timescales, all the way up to the 'macro' scale of the boom–bust speculative bubbles. From a knowledge of these patterns, we tend to form opinions about what will happen next. We would then trade accordingly in an attempt to maximize our wealth. However, the agents we have modelled *cannot* view the past price series in this way: instead they are forced to consider patterns of length m timesteps, where m is a single, fixed number. Patterns of any greater length than this go unnoticed by the agents and hence are not traded upon. This is why these patterns exist in our model. Now, we already know from the concept of arbitrage that when a real pattern is traded upon, it should slowly be removed from the price series. Consider the following pattern:

Time t	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Price $x[t]$	14	10	11	12	13	15	14	10	11	12	13	15	14	10

If we were able to identify this pattern repeating, our best course of action would be to submit a buy order between times $t = 1$ and $t = 2$, that is, $a[1] = 1$. We could then buy the asset at $x[2] = 10$. Between times $t = 5$ and 6, we would submit an order to sell $a[5] = -1$, and sell the asset at $x[6] = 15$. We then continue: $a[7] = 1$, $a[11] = -1$, and so on. This ensures that we always buy at the bottom price and sell at the top. Trading in this way is against the trend since $a[t]D[(t+1)^-] < 0$: it is in effect minority trading. Thus, trading to maximize our wealth with respect to this pattern leads to the weakening of the pattern itself, just as in the Minority Game. We conclude therefore that the *presence* of strong patterns in the MM(W) and other similar market models at low m, is simply due to the *absence* of agents within the model who can identify these patterns and hence arbitrage them out.

From the above discussion, it is clear that in a realistic market model, we should have agents who can analyse the past series of asset price movements over *different* timescales. One possible way of achieving this is to include a heterogeneity in the strategy book pattern length (i.e. memory) m. Within such a framework, agents would look at patterns not only of differing length but also of differing complexity. However, it is more straightforward to propose a generalization to the way the agents interpret the global information of past price movements, in such a way as to allow observation of patterns which occur over different timescales but which have the same complexity (m). This can be easily achieved by allowing the agents to have a natural information bit-length τ such that the global information available to them, $\mu_\tau[t]$, is updated according to the sign of $\Delta x[t+1, t+1-\tau] - L[t+1]$ as opposed to being based simply on $\Delta x[t+1, t] - L[t+1]$. The following example table shows how the above pattern would be encoded for $L[t] = 0$ by agents having $\tau = 1, 2, 3, 4$.

Time t	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Price $x[t]$	14	10	11	12	13	15	14	10	11	12	13	15	14	10
‘Best’ action $a^*[t]$	1	1	1	1	−1	−1	1	1	1	1	−1	−1	1	1
sgn $[\Delta x[t, t-1]]$		−	+	+	+	+	−	−	+	+	+	+	−	−
sgn $[\Delta x[t, t-2]]$			−	+	+	+	+	−	−	+	+	+	+	−
sgn $[\Delta x[t, t-3]]$				−	+	+	+	−	−	−	+	+	+	−
sgn $[\Delta x[t, t-4]]$					−	+	+	−	−	−	−	+	+	−

If each agent only considered the past two bits of information, that is, $m = 2$, then the 'best' strategies for different values of the information bit-length τ would be as shown in the table below. These 'best' strategies have been obtained from inspection of the example table above: specifically, by looking at timesteps where the different $m = 2$ bit-strings $\{--, -+, +-, ++\}$ occur, and then seeing the respective 'best' action $a^*[t]$ given this bit string.

$\mu[t]$	Bit-length τ			
	1	2	3	4
$--$	1	1	1	?
$-+$	1	1	-1	-1
$+-$	1	1	1	1
$++$	?	?	?	1

A question mark next to a particular value of the global information $\mu[t]$ denotes that, for this state, the best action is sometimes $a_{R^*}^{\mu[t]} = 1$ and sometimes $a_{R^*}^{\mu[t]} = -1$. We can, therefore, see that it is only when we include longer timescale patterns $\tau > 2$, that we get a clear signal of when the optimal time to sell occurs ($a_{R^*}^{\mu[t]} = -1$): shorter timeframe patterns give no such clear indication. This then demonstrates that an agent holding strategies of different bit-length τ could identify optimal times to buy and sell, and hence arbitrage patterns of length very much greater than the memory length m.

4.5 Choosing a model

Section 4.4 discussed how a realistic market model could be built from the framework of the El Farol bar problem, by successive adaptation of the model's features to the real-life scenario of financial market trading. This involved not only modifying the strategy book reward structure, which fundamentally changes the style of game, but also changing the subtleties of the agents' composition. We discussed how the inclusion of finite agent resources, and diversity in trading timescales, can yield a model which is free from instabilities and atypical market patterns. There are also subtleties of the market itself which we have not addressed, that is, the presence and actions of a market-maker together with all other possible market-moving influences such as news arrival. However, our original goal was to understand what lies behind some of the more interesting dynamical properties observed in a market. For example, what gives rise to the observed high volatility in today's crowded markets? It will be impossible to answer these types of questions if the model we use to simulate the market is of comparable complexity to the market itself. Instead, we need to focus on a minimal set of underlying assumptions in order to make sense of what is, after all, a very 'complex' system.

In the following sections, we will present and analyse a *minimal* market model in order to address the question concerning the origins of market volatility. We will choose the most basic model we have, which reflects at least *some* major aspects of financial trading and yet also reproduces the stylized facts observed in real financial market data. In particular, we shall investigate the *direct* application of the El Farol bar problem to a financial market. By this we mean that we will preserve the structure of 'complete' strategy books, and the global information given by the past series of price increments relative to $L[t]$. We will then construct a demand for assets using Equation (4.6) and will reward strategy books using Equation (4.5). As discussed earlier, this structure captures the essence of what is important at the time of trading but neglects to take full account of the agents' accumulated positions. Despite this shortcoming, the model may well prove to be an adequate model of trading on short (intra-day) timescales where agents seek to make money on the immediate, per-trade basis.

4.6 The 'El Farol Market Model'

4.6.1 Specifying the model

Let us begin by summarizing how this minimal market model is specified. The model has six parameters:

- N = the number of agents
- m = the 'memory' of the agents
- s = the number of strategy books (i.e. strategies) held by each agent
- r = the minimum score that an agent's best strategy book R^* must have in order for him to participate
- T = the time horizon over which strategy book scores are 'forgotten'
- $L[t]$ = the size of increment against which the asset's movement is judged.

A schematic diagram of the model is shown in Fig. 4.8. The agents all observe a common binary source of information representing recent price movements, of which they only remember the previous m bits. Hence, the global information available to each agent at time t is given by $\mu[t]$, where in decimal notation $\mu[t] \in \{0 \ldots P-1\}$ with $P = 2^m$. Each strategy book $\underline{a}_R$ contains as its elements a_R^{μ}. These elements

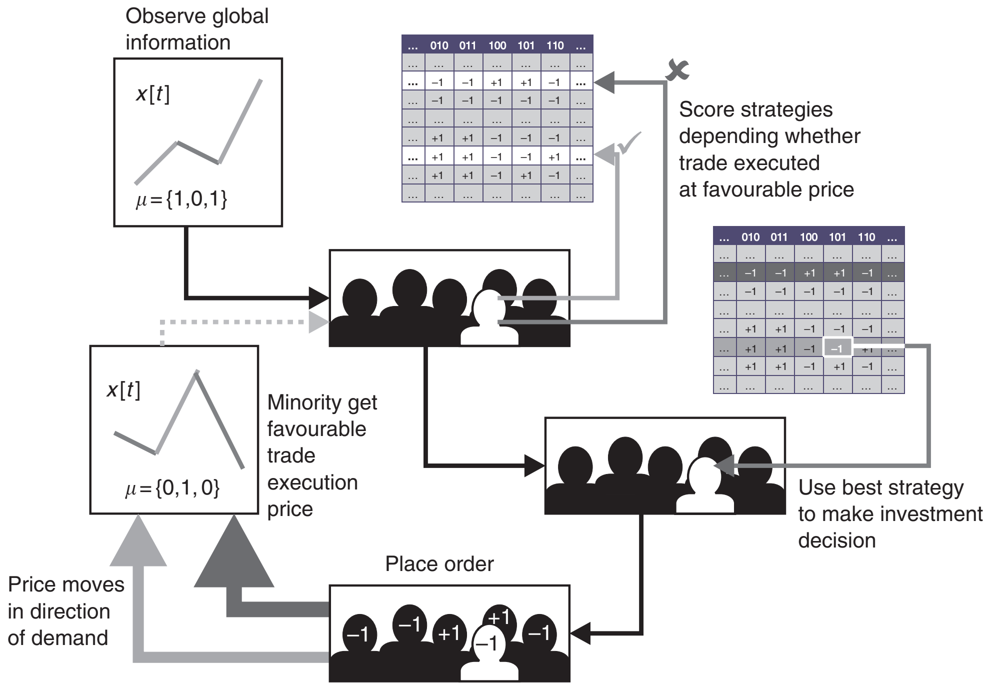

Fig. 4.8 Schematic diagram of the El Farol Market Model, which is a binary multi-agent game. This model represents our minimal market model.

provide an action $\{-1, +1\}$ representing $\{\text{sell}, \text{buy}\}$, for each of the possible P values of the global information μ. There are hence 2^P possible strategy books. The agents are allocated randomly a subset s of these strategy books at the outset of the game: they are not allowed to replace these during the game. The agents keep the score $S_R[t]$ of each of their strategy books' previous successes, regardless of whether the strategy book's investment suggestion was used or not. Following Equation (4.5), success is defined by the reward structure:

$$S_R[t+1] = (1 - 1/T)S_R[t] - a_R^{\mu[t]}\text{sgn}\left[D\left[(t+1)^-\right]/\lambda - L[t+1]\right], \quad (4.12)$$

where the demand for assets immediately prior to the deal execution time $t + 1$, is given by

$$D[(t+1)^-] = n_{\text{buy\_orders}}[t] - n_{\text{sell\_orders}}[t] = \sum_{R=1}^{2^P} n_R[t]a_R^{\mu[t]}. \quad (4.13)$$

Here $n_R[t]$ is the number of agents choosing to follow the investment suggestion of strategy book R at time t. The agents always use their highest performing strategy book R^*, that is, $S_{R^*} = \max[\{S_R\}_s]$, but only participate at a given timestep if this performance has exceeded the confidence-to-trade threshold of r. If $S_{R^*} < r$, then the agents do not participate. The feedback in this market model arises through the updating of the global information $\mu[t]$, where the most recent bit of binary information is defined by the sign of the price movement $\Delta x[t, t-1]$ relative to $L[t]$. Hence, we can calculate the update of the decimal global information variable $\mu[t]$ according to the following expression:

$$\mu[t+1] = 2\mu[t] - PH[\mu[t] - P/2] + H\left[D\left[(t+1)^-\right]/\lambda - L[t+1]\right], \quad (4.14)$$

where $H[x]$ is the Heaviside function. This makes sense because the increment in the asset price is generally an increasing function of the demand, for example

$$x[t] - x[t-1] = D[t^-]/\lambda. \quad (4.15)$$

Despite its unfamiliar appearance, Equation (4.14) for the update of $\mu[t]$ is easy to evaluate. For example, consider the case of $m = 2$ and $L[t] = 0$ for all t. Suppose the recent outcome bit-string is 00 and hence $\mu[t] = 0$: a negative excess demand will give $\mu[t+1] = 2 \cdot 0 - 2^2 \cdot H[0 - (2^2/2)] + 0 = 0$ corresponding to an updated recent outcome bit-string 00, while a positive excess demand will give $\mu[t+1] = 2 \cdot 0 - 2^2 \cdot H[0 - (2^2/2)] + 1 = 1$ corresponding to an updated recent outcome bit-string 01. Now suppose the recent outcome bit-string is 01 and hence $\mu[t] = 1$: a negative excess demand will give $\mu[t+1] = 2 \cdot 1 - 2^2 \cdot H[1 - (2^2/2)] + 0 = 2$ corresponding to an updated recent outcome bit-string 10, while a positive excess demand will give $\mu[t+1] = 2 \cdot 1 - 2^2 \cdot H[1 - (2^2/2)] + 1 = 3$ corresponding to an updated recent outcome bit-string 11.

4.6.2 Parameterizing the El Farol Market Model

We now examine the price $x[t]$ and 'volume' of orders $V[t] = \sum_{R=1}^{2P} n_R[t-1]$. Let us start by deciding on the parameter set $\{N, m, s, r, T, L[t]\}$. The agents in our model represent financial individuals or institutions which are capable of moving the market with their orders. We are thus not considering the actions of small home investors in these models. Consequently, we expect there to be of order $10^2 \rightarrow 10^3$ of these market-moving agents. It thus seems reasonable to set $N = 500$. We have discussed previously that the financial market is likely to be in a 'crowded' phase wherein large groups of agents agree on the same course of action and rush to the market together. In the context of the binary model structure, this necessitates us having many agents holding similar strategy books. This condition requires that $N\ s \gg 2P = 2^{m+1}$ (recall that there are $2P$ strategy books in the RSS). We can satisfy this condition by setting $m = 3$ and $s = 2$. Hence the agents are adaptive because they can change which strategy book they use, and the set of global information states is small, but not trivially so. The parameters r and T together define the number of strategy books which will have a success rate above the confidence-to-trade threshold. If the demand generated by the agents were a zero-mean random variable, we would expect the strategy book reward structure of Equation (4.12) to give success scores $S_R[t]$ with a variance $\sigma^2[S_R]$ given by $\sigma^2[S_R] = 1/(1 - \alpha^2)$, where $\alpha = 1 - 1/T$. If we then said that a reasonable characteristic timescale over which strategy book scores are forgotten was of order $T = 100$ timesteps, this would give $\sigma[S_R] \approx 7$. In reality of course, the demand is not a random variable and we find that the variance of strategy book success scores is smaller than this value due to competition between the strategy books. We thus pick a value of $r = 4$ as a reasonable confidence-to-trade threshold, since it will represent approximately a one-sigma deviation from the random coin-toss success rate of $S_R = 0$. Finally, we note that the parameter $L[t]$—which represents the increment in value of a reference asset (the asset to which the traded asset is being compared)—can in general be time dependent. However, the role of $L[t]$ is simply to shift the reference frame of the model, hence it is not particularly instructive to consider its time dependence here. We do, however, note that if the value of $L[t]$ is changing on a timescale comparable to the timestep size of the model itself, we expect this extra source of dynamical behaviour to be important.[19] For our purposes, we will take $L[t] = L$, that is, the increment in value of the reference asset is a constant.

4.6.3 Reproducing the stylized facts

Having formalized and parameterized the El Farol Market Model, we are ready to run the simulations. Every new simulation of the market will be unique due to the random initial allocation of strategy books among agents. We do not wish here to analyse specific run-dependent anomalies in the output, but will instead focus on general behavioural patterns and statistical features. Specifically we are interested to see whether this model, with the given physically motivated parameterization, reproduces the stylized facts of a financial market. First, let us inspect by eye the model's output.

Figure 4.9 shows the typical evolution of the El Farol Market Model, for different fixed increments in value of the reference asset. For a reference asset which appreciates in value ($L > 0$), the market price of the traded financial asset also appreciates and vice versa. For $L = 0$ we observe an apparently unbiased movement in price for the traded

[19] For example, if $L[t]$ reflects news arrival, and significant news arrives very frequently, then this news-arrival process is likely to be a major factor in determining the dynamics of the market.

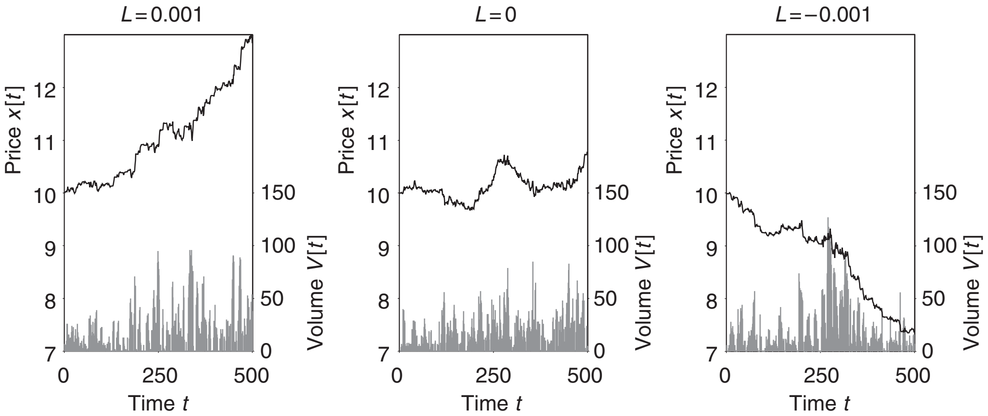

Fig. 4.9 Typical price and volume charts produced by the El Farol Market Model as specified in Section 4.6.1 and parameterized in Section 4.6.2. Each chart represents a different fixed value $L[t] = L$, which corresponds to the constant increment in value of the reference asset.

asset. These results make sense: Equation (4.12) tells us that for $L > 0$, agents tend to reward buying strategies ($a_R^{\mu[t]} = 1$) over selling strategies (and vice versa for $L < 0$). This creates a bias in the demand and consequently the price increment. This mirrors what we observe for financial indices where the reference asset is typically taken to be the rate of interest *and* inflation combined. Historically, during periods of inflation and positive interest, market indices have tended to appreciate in value. The activity pattern of trading in the asset looks from Fig. 4.9 to be interesting and non-random. Activity seems to be clustered, with periods of high volume trading following other similar periods. Large movements in the asset value, compared with the average volatility, also seem to be common. We will now investigate these properties by looking into the statistics of the series of asset price-changes $\Delta x[t+1, t]$ with $L = 0$.

Figure 4.10 shows that the distribution of asset price increments, produced by the El Farol Market Model, deviates significantly from the random case (i.e. Gaussian or so-called 'normal' distribution). The high probability of large movements is reflected in the high kurtosis (peakedness) of the PDF and the deviation of the QQ plot from the diagonal line. Figure 4.11 shows graphically that the linear (i.e. low-order) correlations between the price increments are small and short-ranged: however, the autocorrelation in the local volatility is of much longer range and has large magnitude. The autocorrelation of the local volatility represents a non-linear (i.e. higher order) correlation measure for the price increments, and is practically equivalent to the autocorrelation of absolute returns. Comparing Figs 3.4 and 4.11, we hence see that the El Farol Market Model is capable of capturing both the low-order *and* higher order correlations which are observed in real market data (see Chapter 3).

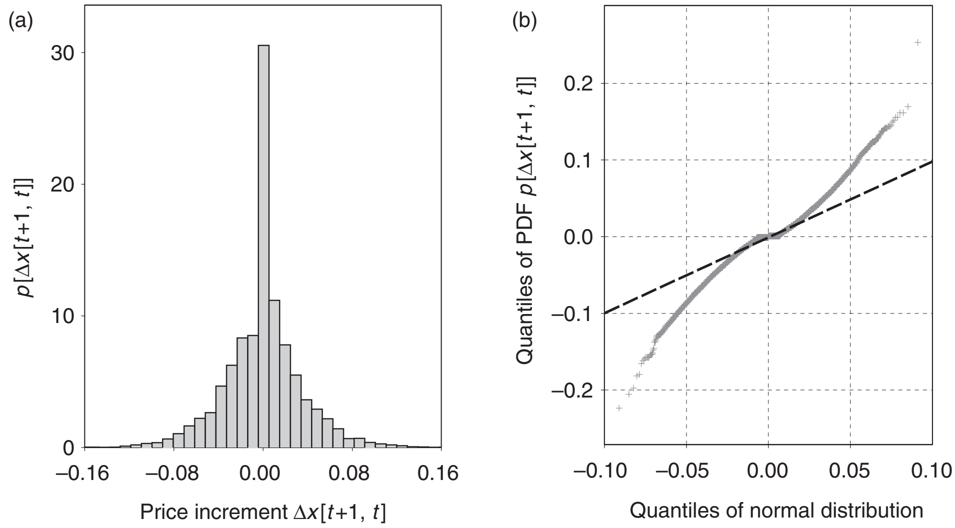

Fig. 4.10 Properties of the series of asset price increments $\Delta x[t+1, t]$ produced by the El Farol Market Model as specified in Section 4.6.1 and parameterized in Section 4.6.2. Here $L = 0$. (a) The histogram represents the probability distribution function (PDF) of the asset price-changes, and exhibits extremely high kurtosis (fat tails). (b) A QQ plot which demonstrates the deviation of this distribution from the random (i.e. Gaussian or so-called 'normal') case. The dashed line is that expected for a Gaussian (i.e. normal) distribution.

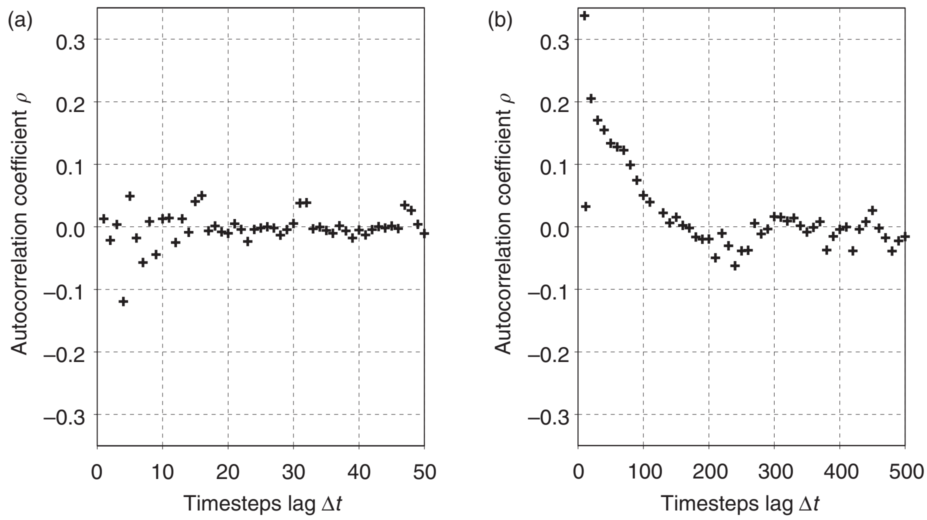

Fig. 4.11 Correlations in the series of asset price increments $\Delta x[t+1, t]$ produced by the El Farol Market Model as specified in Section 4.6.1 and parameterized in Section 4.6.2, with $L = 0$. (a) The autocorrelation of the price increments and (b) the autocorrelation of the local volatility as measured over 10 timesteps.

To summarize, *the statistical features of the El Farol Market Model are consistent with the stylized facts observed in real financial markets*. Of course, this does not mean that the market model we have simulated is either 'correct and unique' or even just 'correct'. However, it does show that the model is capable of exhibiting the kind of rich behaviour that is observed in real financial markets, despite the minimal set of assumptions and parameters.

4.7 Dynamics of the 'El Farol Market Model'

The El Farol Market Model described in Section 4.6, seems to be an easily specifiable system. However, there are subtle choices to be made when writing a computer program to implement the game. These can lead to different results for the time evolution of the game from any two programs, even though the programs are 'identical'. Despite the fact that the market model is exactly as stated earlier, and even though we might use exactly the same parameter set for any two simulations, the numerical results will differ in their specific time evolution because of:

1. Different initial conditions. These initial conditions come in three forms:

   (a) Initial strategy book allocation among the agents. This initial strategy book assignment is random and yet fixed from the outset, hence providing a systematic 'disorder' which is built into each run. In technical terms, different runs will generally have different realizations of this 'quenched disorder'. The strategy book allocation can be represented as a matrix $\underline{\underline{\Omega}}$ for the case $s = 2$ strategies per agent. This $2P \times 2P$ (in the RSS) quenched disorder matrix has as its elements $\Omega_{R,R'}$, where $\Omega_{R,R'}$ is the number of agents who have first been allocated strategy book R and then R'. Since the order of allocation is irrelevant, we can without loss of generality consider a symmetrical matrix $\underline{\underline{\Psi}} = \frac{1}{2}(\underline{\underline{\Omega}} + \underline{\underline{\Omega}}^{\mathrm{T}})$ to describe the quenched disorder in the system. An example matrix is shown below:

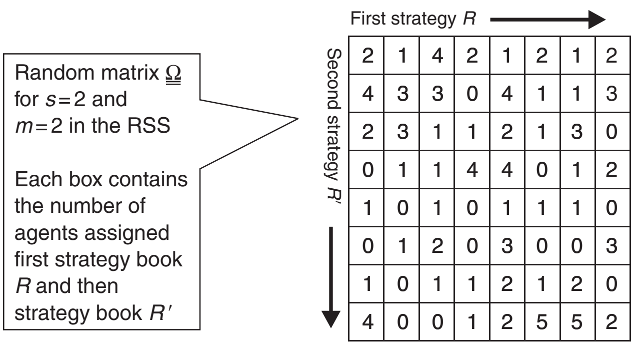

2	1	4	2	1	2	1	2
4	3	3	0	4	1	1	3
2	3	1	1	2	1	3	0
0	1	1	4	4	0	1	2
1	0	1	0	1	1	1	0
0	1	2	0	3	0	0	3
1	0	1	1	2	1	2	0
4	0	0	1	2	5	5	2

(b) Initial success score of the strategy books. If this initial score distribution is not 'typical', then a bias can be introduced into the game which never disappears. In short, the system never recovers from this bias. We will assume that no such initial bias exists, that is, we assume that the initial score-vector distribution is 'typical' and does not introduce any long-term bias. In practice this is achieved, for example, by setting all the initial scores to zero.

(c) Initial history used to seed the simulation. This is not in general an important effect. We assume that any transient effects resulting from the particular history seed will disappear quickly, that is, we assume that the initial history seed does not introduce any long-term bias.

2. Each run has an intrinsic stochasticity because of accidental ties in strategy books' success scores. These ties are typically rare. When they do occur, the rules of the game are such that an unbiased coin-toss will determine which strategy book is used. In the event that the total demand from the agents is equal to λL and hence $\Delta x[t+1,t] = L$, the global information state $\mu[t]$ must also be updated by a coin-toss. Hence, even if two runs have the same initial conditions, as defined above, they will eventually differ in their time evolution due to this stochasticity.

In the rest of this chapter, we will focus on the static (i.e. time-averaged) quantities of the multi-agent game. We note, however, that a fuller treatment which includes the dynamics can also be carried out. The secret to this dynamical description is as follows. The simulation uses coin-tosses to resolve ties in strategy book success scores and to update the global information in the situation where $\Delta x[t+1,t] = L$. These events inject stochasticity into the game's evolution. However, one can average over this stochasticity to yield a description of the game's deterministic dynamics via mapping equations[20]: these mapping equations just involve the vector of all the strategy books' success scores $\underline{S}[t]$, and the global information $\mu[t]$. We note that this can be done generally, *without* making the simplifications necessary to obtain the Minority Game, that is, it can be done for time dependent $L[t]$, general confidence level r, and general time horizon T.

4.8 Statics of the 'El Farol Market Model': the origins of volatility

We will now explore the El Farol Market Model specified in Section 4.6.1 and parameterized in Section 4.6.2. Because of the important role that volatility plays in

[20] For the simpler case of the Minority Game, see Jefferies, P., Hart, M., and Johnson, N. F. (2002) *Phys. Rev. E* **65**, 016105.

the financial world, we have set ourselves the goal of understanding the origins of volatility in our market model, and quantifying this volatility based on the model's parameters. We are aiming at a pedagogical presentation of an analytic derivation of this volatility, hence we will make some further simplifications for the purpose of clarity. Specifically, we set the reference asset to have a zero return, that is, $L[t] = 0$. We will also force all agents to participate in trading at every timestep, that is, the 'volume' $V[t] = N$ for all t. This is achieved by taking the limits $T \to \infty$ and $r \to -\infty$. In this regime, the model we recover is the basic Minority Game of Challet and Zhang.[21] Thus, we have reduced the parameter set to the three parameters $\{N, m, s\}$.

4.8.1 Numerical results for the volatility

The asset price is defined by Equation (4.2). Thus, the volatility of the asset price increments is simply a function of the volatility of the agents' total demand for assets. We denote this volatility in demand simply by $\sigma[D[t^-]] = \sigma$. The volatility σ is then a time-averaged quantity given by the equation

$$\sigma^2 = \frac{1}{n}\sum_{t=1}^{n} D[t^-]^2 - \left(\frac{1}{n}\sum_{t=1}^{n} D[t^-]\right)^2. \tag{4.16}$$

We are thus assuming that the distribution of price increments is stationary over the period $1 \le t \le n$, such that the volatility is a time-independent property of the time-series. This is one reason why we use a simplified model wherein all the agents must participate at each timestep: if the number of active agents was instead a function of time, we would expect that the volatility would become highly time dependent (recall our discussion of volatility clustering in Section 4.6.3). We must also keep this stationarity in mind when we produce a numerical value for σ from the market model simulations: the model takes a while ($n \gg 2^m$ timesteps) to settle into a dynamical steady state, hence it is necessary to discard the initial timesteps of the simulation when making a numerical measurement of the steady-state volatility. For simplicity, we shall also consider the model to have strategy books drawn from the RSS as discussed in Section 4.4.2. It turns out that many of the time-averaged features of the model obtained using the RSS, such as the volatility, are almost identical to those produced by the model using the FSS. The reason for this is that all the important *types* of correlations between strategies (i.e. fully correlated, uncorrelated and anticorrelated)

[21] See note 7.

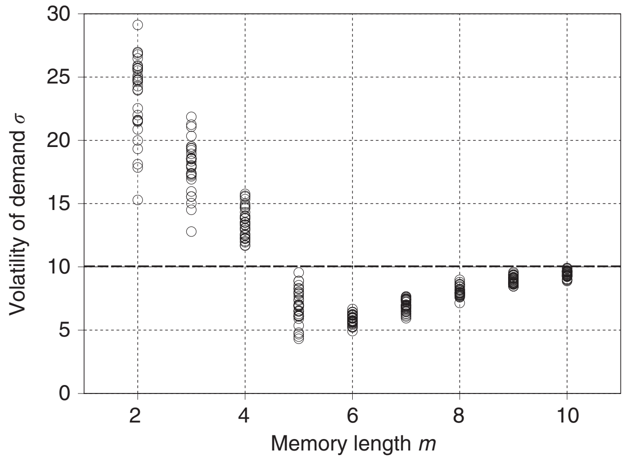

Fig. 4.12 Numerical results for the volatility of total demand σ in a simplified version of the El Farol Market Model (i.e. Minority Game) as a function of 'memory' m. Model parameters are $N = 101$ and $s = 2$. The dashed line at $\sigma = \sqrt{N} = 10.0$ represents the 'random limit', where all the agents simply toss a coin to decide on an investment action.

are present in the RSS. Therefore, the 2^{m+1} strategy books in the RSS provide an adequate representation of the 2^{2^m} strategy books in the FSS.

We now look at some numerical results for this simplifed El Farol Market Model (i.e. Minority Game). Following our earlier discussion, we choose $V[t] = N = 101$ active agents each having $s = 2$ strategy books (i.e. strategies). We then examine the volatility of the model's output as a function of m, the 'memory' or global information bit-string length that the agents use. We run the simulation 32 times for each value of m. Each run of the model corresponds to a different realization of the quenched disorder (i.e. different realization of the disorder in strategy book selection, as discussed in Section 4.7).

Figure 4.12 suggests that there are two distinct 'phases' of the model market, according to the value of m. The low-m phase is characterized by a decrease in σ as m increases: this is the 'crowded' phase we discussed earlier where the number of strategies 2^{m+1} (in the RSS) is small compared to the number of agents N. The high-m phase is characterized by a slow increase in σ towards some limiting value as m increases. This phase is called the 'dilute' phase since the number of strategies 2^{m+1} (in the RSS) is now large compared to the number of agents N.

4.8.2 Qualitative explanation for the variation of volatility

Before looking at a detailed quantitative theory, we first give a simple qualitative picture to explain what is happening. It all relates back to the correlation between strategies, discussed in Section 4.4.2. In the 'crowded' phase, that is, at small m, there will at any one time be a large number of agents who are using the same (e.g. the perceived best) strategy and so will flood into the market as large groups or *crowds*, producing large swings in demand and hence a high volatility as shown. If the memory m of the agents is larger, then the crowd of agents using the same strategy will be smaller simply because many may not hold the best strategy—the chances of a given agent holding the instantaneous best strategy decrease as m increases. There will also be groups of agents who are forced to use the anticorrelated (e.g. the perceived worst) strategy: these can be thought of as *anticrowds* since they cancel out the market action of the crowds at every timestep t *regardless* of the particular history bit-string at that timestep. This cancellation effect causes a reduction in the size of the market volatility. In the dilute phase of very large memory m, it is very unlikely that any two agents will hold the same strategy and so the market can be modelled as a group of independent coin-tossing agents.

Let us make this a bit more quantitative. Consider the oversimplified case of N independent agents each deciding on an investment decision by tossing a coin. Each agent, therefore, provides a random-walk process in terms of increasing or decreasing the demand $D[t^-]$ by 1 asset. Assume for the moment that these coin-tosses are uncorrelated. Then the results of Chapter 2 tell us that the total variance σ^2 for this random walk in excess demand $D[t^-]$, is given by the *sum* of the individual variances produced by each of the N agents. If the agent decides $a^{\mu[t-1]} = 1$, then he contributes 1 to the excess demand $D[t^-]$. If, by contrast, the agent decides $a^{\mu[t-1]} = -1$, then he contributes -1 to the excess demand. In both cases the random-walk 'step-size' is $d = 1$. This coin-tossing agent chooses $a^{\mu[t-1]} = 1$ with probability $p = \frac{1}{2}$, and $a^{\mu[t-1]} = -1$ with probability $q = \frac{1}{2}$. The variance contributed to σ^2 by this agent is therefore given by $4\,pqd^2 = 1$ since $d = 1$ (see *Background Math box*). Summing over all N agents, the total variance in the excess demand σ^2 is given by $4Npqd^2 = N$. Hence the standard deviation (i.e. volatility) of demand is given by $\sigma = \sqrt{N}$ which, for $N = 101$, gives $\sigma = 10.0$ which is the dashed 'coin-toss' line of Fig. 4.12.

In reality, on any given turn of the game, there will be a number of agents using the same, or similar, strategies. Consider the subset of agents n_R using a particular strategy R. Although there is no information available to a given agent about other individual agents, nor is any direct communication allowed between agents, this subset of agents n_R using a particular strategy R will all make the same investment decision at each timestep *irrespective* of the particular history bit-string for that timestep. Hence, they will act as a *crowd*. Since the corresponding random-walk 'step-size' that this crowd contributes is $d \equiv n_R$, this crowd should contribute a variance $4pqd^2 = 4\frac{1}{2}\frac{1}{2}n_R^2 = n_R^2$ to the total variance. However, because of the initial strategy allocation, there may also be a subset of agents $n_{\bar{R}}$ who are using the anticorrelated strategy to R,

Background Math

The expressions below follow directly from the results in Chapter 2, but are summarized here for clarity. Consider a random walk along the y-axis, with step-size $= d$ and number of steps $= N$.

The probability of moving in a positive (negative) direction at each step $= p$, (q), where $p + q = 1$.

The mean displacement $y_{N=1}$ for $N = 1$ is given by:

$$\langle y_{N=1}\rangle = pd + q(-d) = (p - q)d; \quad \text{hence} \quad \langle y_{N=1}\rangle = 0 \text{ if } p = q = \tfrac{1}{2}.$$

To calculate the variance $\sigma^2_{N=1}$ for $N = 1$, we start with $\langle y^2_{N=1}\rangle = pd^2 + q(-d)^2 = (p+q)d^2 = d^2$; hence $\sigma^2_{N=1} \equiv \langle y^2_{N=1}\rangle - \langle y_{N=1}\rangle^2 = d^2 - (p-q)^2 d^2 = d^2[1 - (2p - 1)^2] = 4\,pqd^2$.

For uncorrelated steps, the **variance (or average) of the sum = sum of the variances (or averages)**. The mean displacement y_N for $N \geq 1$ is, therefore, given by:

$$\langle y_N\rangle = N\langle y_{N=1}\rangle = N(p - q)d \quad \text{and hence} \quad \langle y_N\rangle = 0, \text{ if } \ p = q = \tfrac{1}{2}.$$

The variance σ^2_N for $N \geq 1$ is, therefore, given by $\sigma^2_N = N\sigma^2_{N=1} = 4Npqd^2 \equiv \sigma^2$.

Hence $\sigma^2_N \equiv \sigma^2 = Nd^2$, if $p = q = \frac{1}{2}$. Note that $\sigma^2_N \equiv \sigma^2 = N$, if $p = q = \frac{1}{2}$ and $d = 1$.

that is, $\bar{R}$. This second group, the *anticrowd*, makes the *opposite* investment decision to the crowd at each timestep *irrespective* of the particular history bit-string for that timestep. Over the timescale during which these two opposing strategies R and $\bar{R}$ are being played, the fluctuations are determined only by the net crowd-size $n^{\text{eff}}_R = n_R - n_{\bar{R}}$, which constitutes the net step-size of the crowd–anticrowd pair. Hence, the net contribution by this crowd–anticrowd pair to the random-walk variance, is given by $4pqd^2 = \left[n^{\text{eff}}_R\right]^2$. We will now use this result. Suppose strategy R^* is the highest scoring at a particular moment: the anticorrelated strategy $\overline{R^*}$ is, therefore, the lowest scoring at that same moment. In the limit of small m, the size of the strategy space is small. Each agent hence carries a considerable fraction of all possible strategies. Therefore, even if an agent picks $\overline{R^*}$ among his s strategies, he is also likely to have a high scoring strategy. Therefore, many agents will choose to use either R^* itself (if they hold it) or a similar one. Very few agents will have such a poor set of strategies that they are forced to use a strategy similar to $\overline{R^*}$. In this regime, there are practically no anticrowds, and the crowds dominate. Therefore, $n_R \approx N\,\delta_{R,R^*}$ and hence $n_R^{\mathrm{eff}} \approx N\,\delta_{R,R^*}$. Hence, the variance varies as $\sigma^2 \approx N^2$ and is larger than the independent agent limit of N, in agreement with Fig. 4.12. In the limit of large m, the strategy space is very large and agents will have a low chance of holding the same strategy. Even if an agent has several low-scoring strategies, the probability of his best strategy being strictly anticorrelated to another agent's best strategy (hence forming a crowd–anticrowd pair) is small. All the crowds and anticrowds will tend to be of size 0 or 1, implying that the agents act independently. This yields the coin-toss limit discussed above. In the intermediate m region where the minimum in the observed volatility exists, the size of the strategy space is relatively large. Hence some agents may get stuck with s strategies which are all low scoring at a particular timestep. They hence form anticrowds. Considering the extreme case where the crowd and anticrowd are of similar size, we have $n_R^{\mathrm{eff}} \sim 0$ and hence the volatility is essentially zero. This is again consistent with the numerical results. The regime of small volatility will arise for small s since, in this case, the number of strategies available to each agent is small—hence some of the agents may indeed be forced to use a strategy which is little better than the worst-performing strategy $\overline{R^*}$. In other words, the cancellation effect of the crowd and anticrowd becomes most effective in this intermediate m region for small s. Increasing s should make this minimum less marked, as again observed numerically. We can summarize this crowd–anticrowd argument by Fig. 4.13.

4.8.3 Quantitative explanation for the variation of volatility

Consider a given realization of the quenched disorder $\underline{\underline{\Psi}}$. At a timestep t in a given run for this given $\underline{\underline{\Psi}}$, there is a current score-vector $\underline{S}[t]$ and a current history $\mu[t]$, which together define the state of the game. The excess demand $D[(t+1)^-]$ is given by:

$$D[(t+1)^-] \equiv D[\underline{S}[t], \mu[t]]. \tag{4.17}$$

The volatility for a given run corresponds to a *time average* for a given realization of the quenched disorder $\underline{\underline{\Psi}}$ *and* a given set of initial conditions. We will eventually average over many runs, hence effectively average over all realizations of the quenched disorder $\underline{\underline{\Psi}}$ *and* all sets of initial conditions. However, first we focus on a given realization of the quenched disorder $\underline{\underline{\Psi}}$.

We assume that the quantities of interest, that is, the mean and standard deviation of the demand, 'self-average' for the given realization of the quenched disorder $\underline{\underline{\Psi}}$. In other words, we assume that the average over time is equivalent to an average over

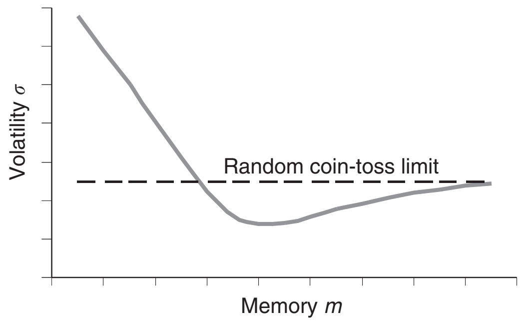

Memory m	$2^{m+1} \ll Ns$	$2^{m+1} \sim Ns$	$2^{m+1} \gg Ns$
Crowd size	Large	Medium	~1
Anticrowd size	Small	Medium	~0
Net crowd–anticrowd pair size	Large $\gg 1$	Small	Small ~1
Number of crowd–anticrowd pairs	$\sim 2^m$ $\ll N$	$\sim 2^m$ $< N$	$< 2^m$ $\sim N$

Fig. 4.13 Diagram explaining the qualitatively different behaviour of the volatility in demand for low, intermediate, and high memory m regimes in terms of crowd and anticrowd sizes.

initial conditions, for the given realization of the quenched disorder $\underline{\underline{\Psi}}$. Such self-averaging in which time averages are taken to be equal to ensemble averages is very common in science, especially physics. For example, when describing the statistical properties of a gas of molecules, there are so many collisions between these molecules that we are guaranteed to explore all the possible microstates of the system either by looking at a time average for a particular system, or an ensemble average over many systems. In the present system, as long as we have (a) 'typical' initial conditions as discussed earlier, plus (b) the stochasticity as a result of coin-tosses to resolve ties, plus (c) the small memory m limit, then the system is indeed self-averaging to a good approximation for a given $\underline{\underline{\Psi}}$. (By 'typical' we mean representative of the set of possible configurations.) Let us rewrite the expression for the demand (Equation (4.13)) using the RSS:

$$D[\underline{S}[t], \mu[t]] = \sum_{R=1}^{2P} n_R^{\underline{S}[t]} a_R^{\mu[t]}. \tag{4.18}$$

In Equation (4.18) above, we have denoted the explicit dependence of $n_R[t]$, which is the number of agents following the investment suggestion of strategy R, on the vector of strategy success scores $\underline{S}[t]$. Now we calculate the average demand, where the average is over time for a given realization of the quenched disorder $\underline{\underline{\Psi}}$. We define $\langle X[t]\rangle_t$ as a time average over the variable $X[t]$ for a given $\underline{\underline{\Psi}}$. By assuming the self-averaging property for a given $\underline{\underline{\Psi}}$, we are essentially assuming that the system is ergodic for a given $\underline{\underline{\Psi}}$: hence all histories are visited with equal frequency in a given run. Hence,

$$\begin{aligned}
\langle D[\underline{S}[t], \mu[t]]\rangle_t &= \sum_{R=1}^{2P} \left\langle a_R^{\mu[t]} n_R^{\underline{S}[t]} \right\rangle_t \\
&= \sum_{R=1}^{2P} \left\langle a_R^{\mu[t]} \right\rangle_t \left\langle n_R^{\underline{S}[t]} \right\rangle_t \\
&= \sum_{R=1}^{2P} \left(\frac{1}{P} \sum_{\mu=0}^{P-1} a_R^{\mu[t]} \right) \left\langle n_R^{\underline{S}[t]} \right\rangle_t \\
&= \sum_{R=1}^{2P} 0 \cdot \left\langle n_R^{\underline{S}[t]} \right\rangle_t = 0. \qquad (4.19)
\end{aligned}$$

Notice that we have averaged over all values of the global information $\mu[t]$ separately, because of our ergodic assumption. We are interested in the fluctuations of the demand about this average value. Hence we move on to consider the volatility of the demand. The variance σ_Ψ of the demand for a particular quenched disorder $\underline{\underline{\Psi}}$, is given from Equations (4.16) and (4.19) by:

$$\begin{aligned}
\sigma_\Psi^2 &= \langle D[\underline{S}[t], \mu[t]]^2\rangle_t - \langle D[\underline{S}[t], \mu[t]]\rangle_t^2 \\
&= \langle D[\underline{S}[t], \mu[t]]^2\rangle_t \\
&= \sum_{R,R'=1}^{2P} \left\langle a_R^{\mu[t]} n_R^{\underline{S}[t]} a_{R'}^{\mu[t]} n_{R'}^{\underline{S}[t]} \right\rangle_t . \qquad (4.20)
\end{aligned}$$

Now we break this double sum into three parts: $\underline{a_R} \cdot \underline{a_{R'}} = P$ (fully correlated), $\underline{a_R} \cdot \underline{a_{R'}} = -P$ (anticorrelated), and $\underline{a_R} \cdot \underline{a_{R'}} = 0$ (uncorrelated). Note that we can only do this decomposition in the RSS. Hence,

$$\begin{aligned}\sigma_\Psi^2 &= \sum_{R=1}^{2P}\left\langle \left(a_R^{\mu[t]}\right)^2 \left(n_R^{S[t]}\right)^2 \right\rangle_t + \sum_{R=1}^{2P}\left\langle a_R^{\mu[t]} a_{\bar{R}}^{\mu[t]} n_R^{S[t]} n_{\bar{R}}^{S[t]}\right\rangle_t \\ &\quad + \sum_{R\neq R'\neq \bar{R}}^{2P}\left\langle a_R^{\mu[t]} a_{R'}^{\mu[t]} n_R^{S[t]} n_{R'}^{S[t]}\right\rangle_t \\ &= \sum_{R=1}^{2P}\left\langle \left(n_R^{S[t]}\right)^2 - n_R^{S[t]} n_{\bar{R}}^{S[t]}\right\rangle_t + \sum_{R\neq R'\neq \bar{R}}^{2P}\left\langle a_R^{\mu[t]} a_{R'}^{\mu[t]}\right\rangle_t \left\langle n_R^{S[t]} n_{R'}^{S[t]}\right\rangle_t \\ &= \sum_{R=1}^{2P}\left\langle \left(n_R^{S[t]}\right)^2 - n_R^{S[t]} n_{\bar{R}}^{S[t]}\right\rangle_t + \sum_{R\neq R'\neq \bar{R}}^{2P}\left(\frac{1}{P}\sum_{\mu=0}^{P-1} a_R^{\mu[t]} a_{R'}^{\mu[t]}\right)\left\langle n_R^{S[t]} n_{R'}^{S[t]}\right\rangle_t \\ &= \sum_{R=1}^{2P}\left\langle \left(n_R^{S[t]}\right)^2 - n_R^{S[t]} n_{\bar{R}}^{S[t]}\right\rangle_t, \end{aligned} \tag{4.21}$$

where strategy $\bar{R}$ is anticorrelated to strategy R. We can write this sum over $2P$ terms as a sum over P terms as follows:

$$\begin{aligned}\sigma_\Psi^2 &= \sum_{R=1}^{2P}\left\langle \left(n_R^{S[t]}\right)^2 - n_R^{S[t]} n_{\bar{R}}^{S[t]}\right\rangle_t \\ &= \sum_{R=1}^{P}\left\langle \left(n_R^{S[t]}\right)^2 - n_R^{S[t]} n_{\bar{R}}^{S[t]} + \left(n_{\bar{R}}^{S[t]}\right)^2 - n_{\bar{R}}^{S[t]} n_R^{S[t]}\right\rangle_t \\ &= \sum_{R=1}^{P}\left\langle \left(n_R^{S[t]}\right)^2 - 2n_R^{S[t]} n_{\bar{R}}^{S[t]} + \left(n_{\bar{R}}^{S[t]}\right)^2\right\rangle_t \\ &= \sum_{R=1}^{P}\left\langle \left(n_R^{S[t]} - n_{\bar{R}}^{S[t]}\right)^2\right\rangle_t. \end{aligned} \tag{4.22}$$

This provides a key general result:

$$\sigma_\Psi^2 = \sum_{R=1}^{P}\left\langle \left(n_R^{S[t]} - n_{\bar{R}}^{S[t]}\right)^2\right\rangle_t.$$

The problem is: how do we work out how many people $n_R^{S[t]}$ are using a given strategy R at a given timestep t? The technique for calculating $n_R^{S[t]}$ will depend greatly on the quenched disorder in the system. A different calculation technique is needed for large, sparsely filled quenched disorder matrices as compared to small densely filled matrices since, for example, the discreteness of agents becomes important in the former case. We will present an analytical treatment of the volatility in these two different regimes.

4.8.3.1 Analytic form for volatility in the crowded regime. We will look at the limiting case where the averaging over the quenched disorder only includes distributions of the matrix $\underline{\underline{\Psi}}$ which are almost flat, that is, we assume that this ensemble averaging is dominated by the matrices $\underline{\underline{\Psi}}$ which are nearly flat. This will be a good approximation for small m since in this limit the standard deviations of the entries $\Psi_{R,R'}$ are much smaller than their means. This result holds because in the crowded low-m regime, there are many more agents than there are available strategies and hence N is greater than the number of entries. The technique for calculating $n_R^{S[t]}$ will involve, in both m-regimes, a re-labelling operation on the strategies. In the crowded regime we will re-label the strategies according to their score. This involves re-writing the sum in Equation (4.22), this time using a strategy label based on a score ranking K as opposed to the decimal form R. Label K is used to denote the rank in terms of strategy score, that is, $K = 1$ is the highest scoring strategy position, $K = 2$ is the second-highest scoring strategy position etc.:

$$S_{K=1} > S_{K=2} > S_{K=3} > S_{K=4} > \cdots \tag{4.23}$$

A given strategy R may at a given timestep have label $K = 1$, while a few timesteps later have label $K = 5$. Because we know that $S_R = -S_{\bar{R}}$ (since all strategy scores start off at zero), we have that $S_K = -S_{\bar{K}}$ where $\bar{K} = 2P + 1 - K$. Rewriting Equation (4.22) gives:

$$\sigma_\Psi^2 = \sum_{K=1}^{P} \left\langle \left(n_K^{S[t]} - n_{\bar{K}}^{S[t]} \right)^2 \right\rangle_t . \tag{4.24}$$

Consider a typical graph of strategy scores as a function of time for a given realization of the quenched disorder $\underline{\underline{\Psi}}$. This graph is represented by Fig. 4.14 (recall we are focusing on the regime of small m).

The ranking (i.e. label) of a given strategy in terms of success score is changing all the time, since the individual strategies have a variation in score which fluctuates rapidly (see e.g. $S_{R=1}[t]$ in Fig. 4.14). This implies that the specific identity of the Kth highest scoring strategy is changing all the time. It also implies that $n_R^{S[t]}$ is changing

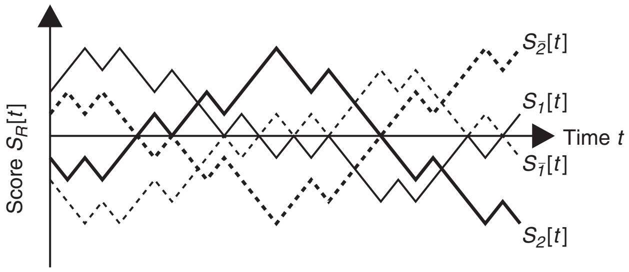

Fig. 4.14 Schematic representation of the variation of strategy success score $S_R[t]$ for two uncorrelated strategies (thick and thin solid lines) and their respective anticorrelated partners (thick and thin dashed lines).

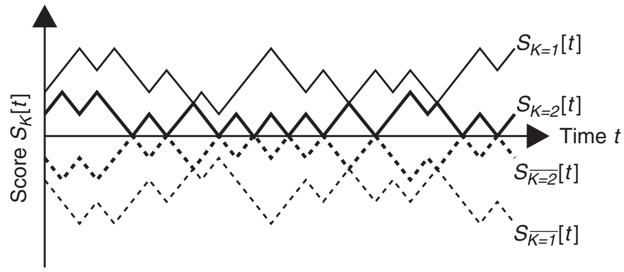

Fig. 4.15 Schematic representation of the strategy success scores from Fig. 4.14, now displayed in terms of $S_K[t]$ where K labels the strategies from highest success score to lowest.

rapidly in time. In order to proceed, we therefore shift our focus to considering the time evolution of the highest scoring strategy, second highest scoring strategy, etc. This has a much smoother time evolution than the time evolution for a given strategy $S_R[t]$. In short, we shift focus from the time evolution of the success score of a given strategy (i.e. from $S_R[t]$) to the time evolution of the success score of the Kth highest scoring strategy (i.e. to $S_K[t]$). From this point of view, Fig. 4.14 is transformed into Fig. 4.15.

We make the assumption that the spread of agents across the strategy space is fairly uniform, that is, $\underline{\underline{\Psi}}$ is a fairly uniform matrix. (This is where we appeal to the model being in the crowded, i.e. low-m, regime.) It therefore makes sense that there will be more agents playing the highest scoring strategy than the second highest scoring etc., that is, it follows from Equation (4.23) and our understanding of the system that we now have:

$$n_{K=1} > n_{K=2} > n_{K=3} > n_{K=4} > \cdots$$

Hence the rankings in terms of highest score and popularity are identical. (Note that the ordering in terms of the labels $\{R\}$ would not be sequential, i.e. it is *not* true that $n_{R=1} > n_{R=2} > n_{R=3} > n_{R=4} > \cdots$.) As shown in Fig. 4.15, the time evolution of the strategy scores is such that the score for a given K tends to fluctuate around a mean value. Hence we expect that the number of traders playing the strategy in position K at any timestep t, will also fluctuate around some mean value. In short, the quantities $n_K^{\underline{S}[t]}$ and $n_{\bar{K}}^{\underline{S}[t]}$ will, therefore, fluctuate in time but far less so than the individual strategy quantities $n_R^{\underline{S}[t]}$ and $n_{\bar{R}}^{\underline{S}[t]}$. Thus, we can make the approximation that:

$$n_K^{\underline{S}[t]} = n_K + \varepsilon_K[t],$$

where $\varepsilon_K[t]$ is assumed to be a white noise term with zero mean and small variance. Here n_K is the mean value. Hence from Equation (4.24), we get:

$$\begin{aligned}
\sigma_\Psi^2 &= \sum_{K=1}^{P} \left\langle (n_K + \varepsilon_K(t) - n_{\bar{K}} - \varepsilon_{\bar{K}}(t))^2 \right\rangle_t \\
&= \sum_{K=1}^{P} \left\langle ((n_K - n_{\bar{K}}) + (\varepsilon_K(t) - \varepsilon_{\bar{K}}(t)))^2 \right\rangle_t \\
&= \sum_{K=1}^{P} \left\langle (n_K - n_{\bar{K}})^2 + (\varepsilon_K(t) - \varepsilon_{\bar{K}}(t))^2 + 2(n_K - n_{\bar{K}})(\varepsilon_K(t) - \varepsilon_{\bar{K}}(t)) \right\rangle_t \\
&\simeq \sum_{K=1}^{P} \left\langle (n_K - n_{\bar{K}})^2 \right\rangle_t = \sum_{K=1}^{P} (n_K - n_{\bar{K}})^2, \qquad (4.25)
\end{aligned}$$

since the latter two terms involving noise will average out to be small. The resulting expression involves no time dependence.

This entire discussion has been for a given realization of the quenched disorder $\underline{\underline{\Psi}}$. We now wish to perform an ensemble average over the various possible realizations of quenched disorder. The values of n_K and $n_{\bar{K}}$ for each K will depend on the precise form of $\underline{\underline{\Psi}}$. Let us denote the ensemble average as $\langle \ldots \rangle_\Psi$, and define for simplicity the notation $\langle \sigma_\Psi^2 \rangle_\Psi = \sigma^2$. We perform this ensemble average on either side of Equation (4.25). Since $\langle (n_K - n_{\bar{K}})^2 \rangle_\Psi$ is just an expectation value of a function of two variables n_K and $n_{\bar{K}}$, we can rewrite it exactly using the joint probability distribution for having n_K *and* $n_{\bar{K}}$, which we call $p[n_K, n_{\bar{K}}]$. Hence:

$$\begin{aligned}
\sigma^2 &= \sum_{K=1}^{P} \left\langle (n_K - n_{\bar{K}})^2 \right\rangle_\Psi \\
&= \sum_{K=1}^{P} \sum_{n_K=0}^{N} \sum_{n_{\bar{K}}=0}^{N} (n_K - n_{\bar{K}})^2 p[n_K, n_{\bar{K}}]. \qquad (4.26)
\end{aligned}$$

So how do we evaluate this? Well, in general it will depend on the joint probability function $p[n_K, n_{\bar{K}}]$ which in turn will depend on the ensemble of quenched disorders $\{\underline{\underline{\Psi}}\}$ which are being averaged over. We will now assume that the probability distribution $p[n_K, n_{\bar{K}}]$ will be sharply peaked around the n_K and $n_{\bar{K}}$ values given by the expected values for a flat quenched-disorder matrix $\underline{\underline{\Psi}}$. Let us call these values n_K^{flat} and $n_{\bar{K}}^{\text{flat}}$. Hence $p[n_K, n_{\bar{K}}] = \delta_{n_K, n_K^{\text{flat}}} \delta_{n_{\bar{K}}, n_{\bar{K}}^{\text{flat}}}$ and so:

$$\sigma^2 = \sum_{K=1}^{P} \sum_{n_K=0}^{N} \sum_{n_{\bar{K}}=0}^{N} (n_K - n_{\bar{K}})^2 p[n_K, n_{\bar{K}}] \simeq \sum_{K=1}^{P} \left(n_K^{\text{flat}} - n_{\bar{K}}^{\text{flat}} \right)^2 .$$

This represents our main result for the crowded regime:

$$\sigma^2 = \sum_{K=1}^{P} \left(n_K^{\text{flat}} - n_{\bar{K}}^{\text{flat}} \right)^2 . \tag{4.27}$$

We now calculate explicit expressions for the case of a flat quenched disorder matrix $\underline{\underline{\Psi}}$. Each entry of the matrix will have a mean of $N(1/2P)^s$ agents, for general s. We can use this to then calculate the expected number of agents who are playing strategy K, that is, n_K^{flat}. In particular, for $s = 2$, we can count the number of agents n_K^{flat} by adding up elements in a reshuffled version of the (flat) quenched disorder matrix, $\Psi_{K,K'}$. Agents will only play strategy $K = \kappa$ if they do not hold any better strategy $K < \kappa$. Thus, the contributing elements of $\Psi_{K,K'}$ to n_K^{flat} are those with $\{K, K'\} = \{\kappa, K \geq \kappa\}$ and $\{K, K'\} = \{K \geq \kappa, \kappa\}$. This is illustrated in the diagram below for $m = 2$:

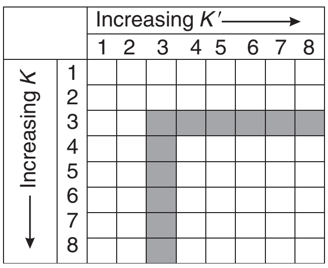

Consider, for example, the expected number of agents using the strategy occupying position $K = 3$ in the list ordered by strategy scores, for $m = 2$. Any agent *using* the strategy in position $K = 3$ cannot have any strategy with a higher position, by definition of the rules of the game (the agents use their highest scoring strategy). Hence, the agents using the strategy in position $K = 3$ are represented by the shaded bins. Since we are assuming that the occupation of the bins is uniform ($\Psi_{K,K'} = N(1/2P)^{s=2}$), the expected number of agents using the strategy in position $K = 3$ is given by:

$$n_K^{\text{flat}} = N\left(\frac{1}{8}\right)^2 \sum_{\text{filled elements}} \Psi_{K,K'} = N\left(\frac{1}{64}\right)((8-3)+(8-3)+1) = \frac{11N}{64}.$$

For more general m and K values, this becomes:

$$n_K^{\text{flat}} = N\left(\frac{1}{2P}\right)^2 ((2P-K)+(2P-K)+1) \equiv \frac{(2^{m+2}-2K+1)N}{2^{2(m+1)}}, \quad (4.28)$$

where we have used $P \equiv 2^m$. Likewise, we have that

$$\begin{aligned} n_{\bar{K}}^{\text{flat}} &= N\left(\frac{1}{2P}\right)^2 ((2P-\bar{K})+(2P-\bar{K})+1) \equiv \frac{(2^{m+2}-2\bar{K}+1)N}{2^{2(m+1)}} \\ &= \frac{(2^{m+2}-2(2^{m+1}-K+1)+1)N}{2^{2(m+1)}} = \frac{(2K-1)N}{2^{2(m+1)}}, \end{aligned}$$

where we have used the identity $\bar{K} = 2P + 1 - K \equiv 2^{m+1} + 1 - K$. We emphasize that the above results depend on the assumption that the averages are dominated by the effects of flat distributions for the quenched disorder matrix $\underline{\underline{\Psi}}$, and hence will only be quantitatively valid for low m. Using Equation (4.27), we therefore have that:

$$\begin{aligned} \sigma^2 &= \sum_{K=1}^{P} \left(n_K^{\text{flat}} - n_{\bar{K}}^{\text{flat}}\right)^2 \\ &= \sum_{K=1}^{P} \left(\frac{(2^{m+2}-2K+1)N}{2^{2(m+1)}} - \frac{(2K-1)N}{2^{2(m+1)}}\right)^2 = \frac{N^2}{2^{4m+2}} \sum_{K=1}^{P} (2^{m+1}-2K+1)^2. \end{aligned}$$

Expanding out the bracket and using the standard relations:

$$\sum_{r=1}^{n} 1 = n, \quad \sum_{r=1}^{n} r = \frac{1}{2}n(n+1), \quad \sum_{r=1}^{n} r^2 = \frac{1}{6}n(n+1)(2n+1),$$

we obtain:

$$\sigma^2 = \frac{N^2}{2^{4m+2}} \sum_{K=1}^{P} (2^{m+1}-2K+1)^2 = \frac{N^2}{3 \cdot 2^m}(1 - 2^{-2(m+1)}).$$

Hence

$$\sigma = \frac{N}{\sqrt{3}\,2^{m/2}}(1 - 2^{-2(m+1)})^{1/2}, \tag{4.29}$$

which should be valid for small m. Figure 4.16 shows the analytical result of Equation (4.29) together with the numerical results for the volatility for different realizations of the quenched disorder $\underline{\underline{\Psi}}$. It can be seen that with our simple treatment of the system, we have managed to capture the dependence of the volatility on

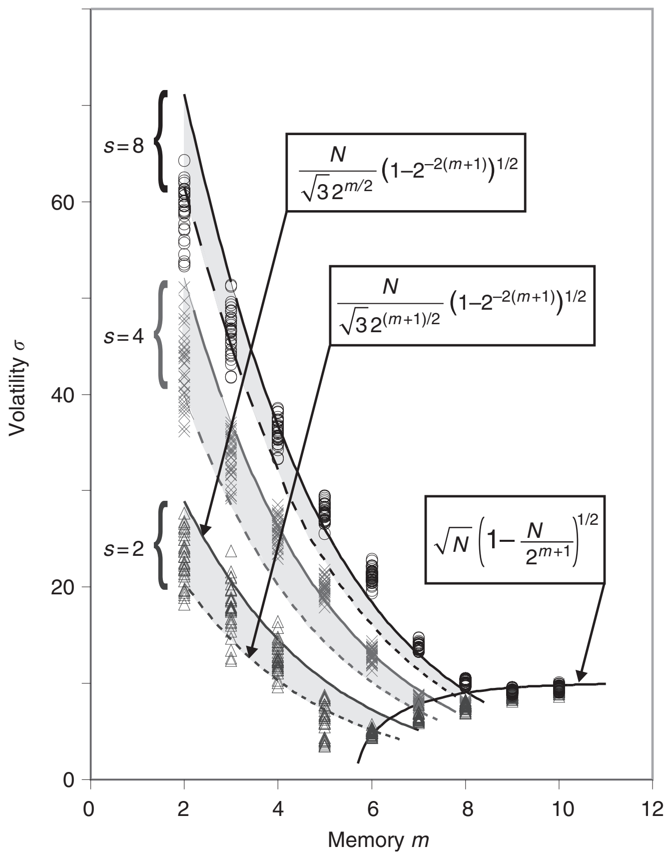

Fig. 4.16 Analytical approximations to the volatility in demand for a simplified version of the El Farol Market Model (known as the Minority Game) as a function of agent memory m. Symbols represent numerical results. Plot corresponds to $N = 101$ active agents. Displayed are results for $s = 2, 4$, and 8 strategies per agent. Separate analytical approximations are made for the crowded (low-m) and dilute (high-m) regimes. In the crowded regime, the upper/lower analytical estimates (solid/dashed lines) represent entirely ordered/disordered rankings of strategies. (Entirely ordered would be $n_{K=1} > n_{K=2} > \cdots n_{K=2P}$).

the model parameters in the low-m regime. We note that the analytical approximation of Equation (4.29) gives a volatility which is slightly higher than the numerical ensemble-averaged result. The main reason for this is the effect of non-zero disorder in $\underline{\underline{\Psi}}$ (i.e. non-flat matrices). Whilst our calculation here assumed an ensemble of quenched disorder matrices $\{\underline{\underline{\Psi}}\}$ which were essentially flat, this will only be the case in practice for $N \to \infty$. As will be shown in Section 4.8.3.2, the effect of finite disorder will be to destroy the ordered ranking of strategies $n_{K=1} > n_{K=2} > \cdots n_{K=2P}$. Without this clear ranking of strategies, the assumptions leading to Equation (4.29) are inappropriate. Section 4.8.3.2 shows that an alternative form of Equation (4.26) (i.e. Equation (4.34)) can then be used. Using Equation (4.34) and assuming no strategy ordering (i.e. $f_{Q',\bar{Q}} = 1/(2P)$) we then arrive at:

$$\sigma = \frac{N}{\sqrt{3}\, 2^{(m+1)/2}}(1 - 2^{-2(m+1)})^{1/2}. \tag{4.30}$$

Equations (4.29) and (4.30) shown in Fig. 4.16 thus represent upper and lower analytical estimates for the volatility corresponding to entirely ordered and entirely disordered strategy rankings, respectively. In general, one would expect that the numerical results would lie within these bounds; this is borne out in Fig. 4.16. Note that the low-m analytic curves can be obtained for any $s \geq 2$ using the generalized form of Equation (4.28):

$$n_K^{\text{flat}} = N(1/2P)^s[(2P - K + 1)^s - (2P - K)^s]. \tag{4.31}$$

We show the analytical approximation and numerical results for $s = 2, 4$, and 8 strategies per agent in Fig. 4.16.

Equation (4.27) has a very simple interpretation in terms of the crowd–anticrowd idea presented in Section 4.8.2, since it represents the sum of the variances for each crowd–anticrowd pair. For a given strategy K, there is an anticorrelated strategy $\bar{K}$. The n_K agents using strategy K are doing the opposite to the $n_{\bar{K}}$ agents using strategy $\bar{K}$ irrespective of the global information bit-string. Hence the average effective group-size for each crowd–anticrowd pair is $n_K^{\text{eff}} = n_K^{\text{flat}} - n_{\bar{K}}^{\text{flat}}$: this represents the net step-size of this crowd–anticrowd pair's random walk. Hence the net contribution by this crowd–anticrowd pair to the variance is given by

$$(\sigma^2)_{K\bar{K}} = 4pqd^2 = \left(n_K^{\text{eff}}\right)^2 = \left(n_K^{\text{flat}} - n_{\bar{K}}^{\text{flat}}\right)^2. \tag{4.32}$$

Since these crowd–anticrowd pairs incorporate all the strong correlations, we can safely assume that the separate crowd–anticrowd pairs execute random walks which are uncorrelated with respect to each other.[22] Hence, the total variance is given by the sum of the individual variances:

$$\sigma^2 = \sum_{K=1}^{P} (\sigma^2)_{K\bar{K}} = \sum_{K=1}^{P} \left(n_K^{\text{flat}} - n_{\bar{K}}^{\text{flat}}\right)^2$$

in agreement with Equation (4.27).

4.8.3.2 Analytic form for the volatility in the dilute regime. If the ensemble of quenched disorder matrices $\{\underline{\underline{\Psi}}\}$ contains a large proportion of non-flat member matrices,[23] we have to approach the averaging over this ensemble of quenched disorders in a different manner than in Section 4.8.3.1. The reason is that the PDF of the number of agents using the Kth most successful strategy, is no longer well approximated by $p[n_K, n_{\bar{K}}] = \delta_{n_K, n_K^{\text{flat}}} \delta_{n_{\bar{K}}, n_{\bar{K}}^{\text{flat}}}$. As m increases, the quenched disorder matrix $\underline{\underline{\Psi}}$ will become increasingly sparse. All the (integer/half-integer) entries $\Psi_{R,R'}$ must add up to give N, the total number of agents: hence as the number of entries $2^{m+1} \times 2^{m+1}$ grows, the matrix becomes filled with an increasing number of zeros. The number of agents using each strategy then tends towards either being zero or one. Hence, $n_R[t]$ is dominated by the disorder rather than the success of the strategy. Furthermore, due to the minority nature of the model, strategies being used will tend to perform less well than strategies which are not used. Thus, strategies which are held by more agents will be less successful than those which are held by few (or no) agents. This leads to a 'market impact force' on n_K. In the low-m crowded regime, the more successful a strategy was, the more agents would play it. The market impact effect was less important there since the number of agents holding each strategy was roughly equal (flat $\underline{\underline{\Psi}}$). In the high-$m$ dilute regime, the opposite is true and the market impact dominates. Hence, the more agents that hold a given strategy, the less successful that strategy is on average. This effect can be seen clearly in Fig. 4.17.

Figure 4.17 shows that the form $p[n_K, n_{\bar{K}}] = \delta_{n_K, n_K^{\text{flat}}} \delta_{n_{\bar{K}}, n_{\bar{K}}^{\text{flat}}}$ calculated as in Section 4.8.3.1 neglecting market impact, is not a good approximation for the high-m regime. In this case, the general analysis is more complicated, and should incorporate the dynamical 'market impact' behaviour of the model. However, we will now develop an approximate theory along slightly different lines which gives good agreement with the numerical results.

22 In the RSS, strategies are either completely anticorrelated or completely uncorrelated to each other.

23 The term non-flat quenched disorder matrix $\underline{\underline{\Psi}}$ refers to a case where the standard deviation of each entry $\Psi_{R,R'}$ is large in comparison with the mean value of the entry. This will become increasingly likely as m increases.

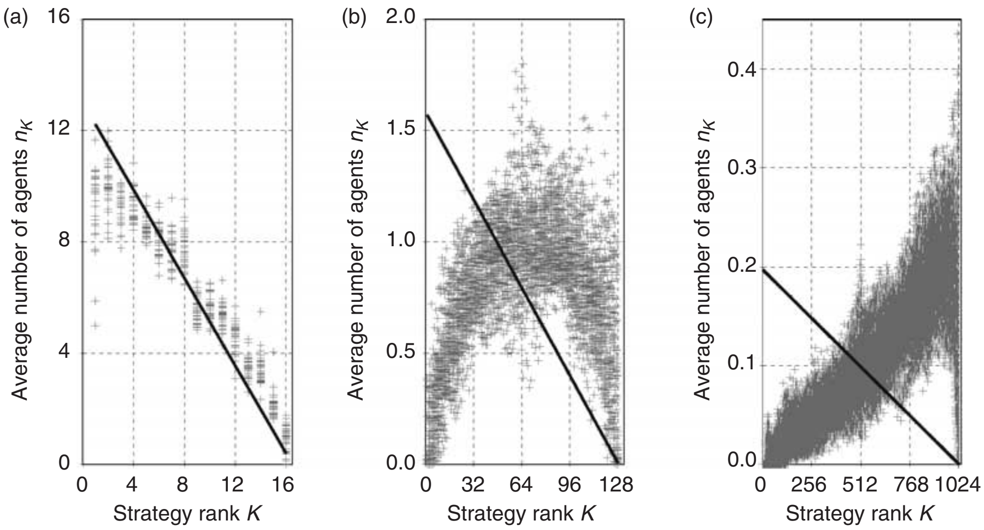

Fig. 4.17 Average number of agents n_K using the Kth most successful strategy book. The crosses represent time averages for 32 different runs. The model parameters were $N = 101$ and $s = 2$ with the memory m equal to (a) 3, (b) 6, and (c) 9. The lines correspond to n_K^{flat} calculated for the low-m regime as described in Section 4.8.3.1.

For the case of ensembles containing a significant number of non-flat quenched disorder matrices $\underline{\underline{\Psi}}$, we can make the following statements:

1. By definition of the labels $\{K\}$, we retain the ranking in terms of success score, that is, it is always true that $S_{K=1} > S_{K=2} > S_{K=3} > S_{K=4} > \cdots$
2. However, the disorder in the matrix $\underline{\underline{\Psi}}$ distorts the number of traders playing a given strategy away from the flat-matrix results. Hence, we *do not* in general have $n_{K=1} > n_{K=2} > n_{K=3} > n_{K=4} > \cdots$

Hence the rankings in terms of highest success score and popularity are no longer identical. In general, we have instead:

$$n_{K'} > n_{K''} > n_{K'''} > \cdots$$

where the label $K' \neq 1$, $K'' \neq 2, \ldots$. However, we can introduce a new label $\{Q\}$ which will rank the strategies in terms of popularity, that is,

$$n_{Q=1} > n_{Q=2} > n_{Q=3} > \cdots$$

where $Q = 1$ represents K', $Q = 2$ represents K'', etc. We now return to the original general form for the volatility, averaged over time, initial conditions and quenched disorder (Equation (4.26)), and rewrite it slightly:

$$\begin{aligned}\sigma^2 &= \sum_{K=1}^{P} \sum_{n_K=0}^{N} \sum_{n_{\bar{K}}=0}^{N} (n_K - n_{\bar{K}})^2 p[n_K, n_{\bar{K}}] \\ &= \frac{1}{2} \sum_{K=1}^{2P} \sum_{n_K=0}^{N} \sum_{n_{\bar{K}}=0}^{N} (n_K - n_{\bar{K}})^2 p[n_K, n_{\bar{K}}] \\ &= \frac{1}{2} \sum_{K=1}^{2P} \sum_{K'=1}^{2P} \left(\sum_{n_K=0}^{N} \sum_{n_{K'}=0}^{N} (n_K - n_{K'})^2 p[n_K, n_{K'}] \right) f_{K',\bar{K}}, \end{aligned} \tag{4.33}$$

where $f_{K',\bar{K}}$ is the probability that K' is the anticorrelated strategy to K (i.e. $\bar{K}$) and is hence given by $f_{K',\bar{K}} = \delta_{K',2P+1-K}$. This manipulation is exact so far. We now switch to the popularity labels $\{Q\}$ to give:

$$\sigma^2 = \frac{1}{2} \sum_{Q=1}^{2P} \sum_{Q'=1}^{2P} \left(\sum_{n_Q=0}^{N} \sum_{n_{Q'}=0}^{N} (n_Q - n_{Q'})^2 p[n_Q, n_{Q'}] \right) f_{Q',\bar{Q}}. \tag{4.34}$$

Consider any particular strategy which was labelled previously by K and is now labelled by Q. Unlike in our treatment of the flat disorder matrix, we are *not* guaranteed that strategy Qs anticorrelated partner $\bar{Q}$ will lie in position $\bar{Q} = 2P + 1 - Q$. This is because of the relabelling operation: all we can say is that the strategy R has changed label from $K \to Q[K]$ while the anticorrelated strategy $\bar{R}$ has changed label from $\bar{K} \to \bar{Q}(\bar{K})$ and that in general $\bar{Q} \neq 2P + 1 - Q$. We thus have two tasks: first to deduce the probability distribution $p[n_Q, n_{Q'}]$ and second to deduce the probability distribution $f_{Q',\bar{Q}}$. In the high-m regime, the number of available strategies is so much larger than the number of agents, that the number using a given strategy n_R is very likely to be zero, slightly likely to be one and almost never anything greater than one. Thus, when the strategies are ranked in terms of the number of agents using them, we have to a good approximation that $n_{Q=1}, n_{Q=2}, \ldots, n_{Q=N} = 1$ and $n_{Q=N+1}, \ldots, n_{Q=2P} = 0$. This gives us:

$$p[n_Q, n_{Q'}] \simeq \delta_{n_Q, H[N-Q]} \delta_{n_{Q'}, H[N-Q']}. \tag{4.35}$$

Next we turn our attention to $f_{Q',\bar{Q}}$, the probability that the strategy Q' is anticorrelated to the strategy Q. First we note that within the groups $n_{Q=1} \ldots n_{Q=N}$ and $n_{Q=N+1} \ldots n_{Q=2P}$, the number of agents using each strategy is the same, consequently the strategies' rankings within these groups are arbitrary. Second we reiterate that the disorder among strategies, rather than the success, is the main factor in determining whether they are used by agents or not. This implies that there is no particularly strong reason why a strategy in group $n_{Q=1} \ldots n_{Q=N}$ should have its anticorrelated partner in either the same group or group $n_{Q=N+1} \ldots n_{Q=2P}$. These two facts combined, lead us to the conclusion that a good approximation to $f_{Q',\bar{Q}}$ would be to say that there is roughly equal probability of finding the anticorrelated partner of Q *anywhere* in the range $Q' = 1 \ldots 2P$, that is, we have:

$$f_{Q',\bar{Q}} \simeq 1/(2P) \tag{4.36}$$

for all Q', and all Q. Now we can substitute the approximate forms from Equations (4.35) and (4.36) into Equation (4.34). This then gives:

$$\sigma^2 = \frac{1}{4P} \sum_{Q=1}^{2P} \sum_{Q'=1}^{2P} (H[N - Q] - H[N - Q'])^2. \tag{4.37}$$

The value of each summand term in Equation (4.37) is either zero or one. If we count the number of times the summand is equal to one, we get:

$$\sigma^2 = (1/4P)2N(2P - N).$$

Hence, we have the following approximate form for the volatility in the high-m regime, which is valid for any s:

$$\sigma = \sqrt{N} \left(1 - \frac{N}{2^{m+1}}\right)^{1/2}. \tag{4.38}$$

Figure 4.16 shows the analytical approximation in Equation (4.38) for the volatility in the high-m regime. Again we seem to have captured the dependence of the volatility on the model parameters. Just as we observed in the numerical simulations, the volatility tends towards the random, coin-toss limit of $\sigma = \sqrt{N}$ as the memory $m \to \infty$. Again we see that our analytical approximation of Equation (4.38) is a little higher than the numerical results. This time the culprit is our analytical approximation for $f_{Q',\bar{Q}}$. It is quite simple to derive a slightly more sophisticated approximation for $f_{Q',\bar{Q}}$ which incorporates the agents' behaviour in picking their most successful strategy. For general s, this then yields:

$$\sigma = \sqrt{N} \left(1 - \frac{Ns - 1}{2^{m+1}}\right)^{1/2}, \tag{4.39}$$

which gives a better fit to the observed ensemble-averaged volatility.

To summarize, our analytical approach has managed to capture the essential interactions driving the variation in volatility for a simplified version of the El Farol Market Model known as the Minority Game. The present analytic results can be extended to account for different model variations. An example is that of 'stochastic strategy picking'. In this version, agents have $s = 2$ strategies each but choose to use their most successful strategy with a probability less than 1, that is, they have a finite probability of using their worst-scoring strategy. The effect of the stochastic strategy picking is to break up the crowds, hence reducing the volatility in the low-m regime.[24] Similarly, one can treat a population of agents which contains a sub-population of stochastic strategy pickers, while the remainder choose their best strategy with probability 1. More generally these analytic approaches to understanding agent-based models based on crowd–anticrowd interaction, need not be limited to finance-related phenomena. Indeed there is currently much work on agent-based systems in computing science, artificial intelligence studies, and in the biological and social sciences.[25] There is even cause to believe that such models and methodologies may help with the more abstract, science-related Holy Grail of understanding complex systems as a whole. Only time will tell.

[24] For discussions and analytic models which we have presented concerning stochastic strategy picking, see Jefferies, P., Hart, M., Johnson, N. F., and Hui, P. M. (2000) *J. Phys. A: Math. Gen.* **33**, L409; Hart, M., Jefferies, P., Johnson, N. F., and Hui, P. M. (2001) *Phys. Rev. E* **63**, 017102. See Johnson, N. F., Hui, P. M., Jonson, R., Lo, T. S. (1999) *Phys. Rev. Lett.* **82**, 3360, for a stochastic model which demonstrates explicitly the spontaneous formation of crowds–anticrowds in a competing population.

[25] See the work of D. Wolpert and K. Tumer at NASA Ames Research Center, www.nasa.arc.gov, on Collectives.