
2.1 Introduction

The aim of this chapter will be to introduce the reader to models of financial markets inspired by psychology. The chapter can in a sense be seen as a prelude to the rest of the book, but with the caveat that it only treats individual, or averaged individual, behavior without taking into account collective (sociological) effects on price formation in the financial markets. Collective effects and complexity models describing financial market price formation will then be introduced in the following chapters.

The origins of behavioral finance can be traced back to questions about the validity of the assumption of rational expectations in decision-making and the theory based upon it, which surfaced in the 1980s. Criticisms were notably raised by a succession of discoveries of anomalies, and in particular the evidence of excess volatility of returns described in Sect. 1.2. Since the reported excess volatility raised some of the very first empirically founded questions relating to the efficient market theory, it became a subject of heated and continuing academic dispute, but the details of these controversies will not be presented to the reader in this book.

The field of decision-making itself became a research topic in the field of psychology in the 1950s, through the work of Ward Edwards, and also Herbert A. Simon [127], who introduced the concept of decision-making based on *bounded* rationality. However, it was not until the work of Daniel Kahneman and Amos Tversky that results from cognitive psychology found their way into economics and finance. In this chapter, we shall give an overview of the kind of irrational beliefs that commonly enter into decision-making. We shall also discuss the implications for financial markets of the Nobel Prize winning prospect theory due to Kahneman and Tversky.

After introducing the reader to various cognitive processes and human biases, we return to the general problem of the way human emotional and cognitive biases impact pricing in financial markets. One issue in particular is, if such human biases occur in the context of pricing in financial markets, how will they manifest

themselves and how can one detect this? Typically, biases are put forward as postulates in the field of behavioral finance, and have been tested in well-controlled laboratory experiments. However, the impact they might have in financial markets is still very much disputed. Critics typically argue that experimentally observed behavior is not applicable to market situations, since learning and competition will ensure at least a close approximation of rational behavior [46]. To probe what could be the impact and signature of behavioral biases in financial markets, it is therefore important to suggest tests and tools that apply directly to financial market data.

2.2 Cognitive Processes: The Individual Level

Almost everyone would agree that decisions made by humans differ from the predictions of rational models. Opinions differ, however, as to whether these differences can be treated as noise that will cancel out in a higher level of description, or whether the difference is so fundamental that standard economic theory cannot be treated as an adequate model of human economic behavior. The difference exists on both the individual and the social level. The line of research in cognitive psychology that resulted in prospect theory, which we will explain shortly, clearly pointed out how individual decision-makers differed from the rational model.

Individual decisions and judgments differ from the prescriptive models in many ways. These differences involve cognitive processes, motivation, emotion, self-structure, and personality factors. The decisions of individuals also strongly depend on how the alternatives are framed, rather than just being based on objective probabilities and outcomes. Individuals use heuristics to estimate probabilities rather than deriving them from the frequencies of occurrences of events. Availability heuristics, for example, describes a process where the decision-maker estimates the probability of events on the basis of how easy it is to recall those events from memory.

Other challenges to the view that humans are rational decision-makers come from cognitive psychology. The results of many experiments show that humans do not use the rules of mathematics and logic as suggested by decision theory, but instead base their judgments on schemas. This leads to a dependence of information processing on context. Child street vendors in Brazil, for example, are perfectly capable of performing arithmetic operations using money in the context of selling, but are unable to repeat the same operations on abstract numbers in a laboratory situation [23]. Wason's card selection task provides another example of this [70].

Imagine that you are working in a company producing games. Your task is to check whether the cards follow the rules of the game. For a certain game that the company produces all the cards must follow the rule: If a card has a vowel on one side, then it must have an even number on the other side. You know that all the cards have a number marked on one side and a letter on the other. In front of you there are four cards. Marked on their upper sides are the following:

E K 5 4

Which cards do you have to turn over to make sure that none of them violates the rule? The correct answer is 'E' and '5'. This is because, if there were an odd number, say 3, on the other side of E (which is a vowel), this card would violate the rule. But the rule would also be violated if there were a vowel, say A, on the other side of 5, since the rule dictates that a vowel on one side implies an even number on the other. Nothing on the other side of K would violate the rule, since the rule does not state what should be on the other side of consonants. Using the same reasoning the card with 4 could not possibly violate the rule. This puzzle is quite difficult, and even people trained in mathematics often have problems finding the correct answer. In the original study less than 10 % of individuals correctly solved the puzzle. The vast majority of people make mistakes in this puzzle [147]. This is because people do not usually use the rules of formal logic in their reasoning, combined with the fact that this example does not look similar to any problems they have encountered so far.

Now consider the following puzzle. Imagine you work in a bar. The rule is that, if a customer drinks alcohol, that customer must be over 18 years old. The bar is dark and you see someone's hand holding a glass of whiskey and another's holding a Coca Cola, but you do not clearly see the faces associated with the hands. You also see two other individuals, one with the face of someone who looks about 16 years old and another with a face that looks over 40 years old. Which of the four individuals do you need to check in order to know how old they are or what they are drinking? The answer is straightforward and usually people have no problems with an answer: one should check the age of the person drinking whiskey and determine what the 16 year old is drinking. Logically, the card game and the bar example are equivalent. The reason why the bar example is so much easier is that it follows a permission schema, a schema that is often used in reasoning in everyday life.

The conclusion we can draw is that, even in problems that have clear formal solutions, people usually use schemas. So to solve problems and make decisions, people are often likely to search for the most relevant schema and use that to arrive at the answer. It follows that, if more than one schema is available, the choice of schema is likely to influence decisions. This effect is often called framing [141]. Therefore, the context in which the decision is made can influence the choice of schema, and the chosen schema will be used as the tool for decision-making.

Another set of challenges to the view of humans as rational decision-makers comes from personality psychology and social psychology. Human decision-making is shaped by various motives. It also depends on emotions.

2.2.1 Motives

In economics it is assumed that humans maximize their outcomes. These outcomes are usually defined in terms of the expected value that would result from the choice. When making a decision, however, individuals often try to maximize more than one outcome. For example, when making money, individuals may also be trying to win approval or make new friends. Choices that seem irrational from the perspective

of one goal may turn out to be rational if we take all the goals into consideration. It is also the case that decisions that satisfy all the goals of an individual will be preferred and made faster than decisions that satisfy only some goals but frustrate other goals [78].

Economic theory assumes self-interest. Psychological research has revealed that people's choices may be governed by different motives and value orientations [61, 97]. Although prevalence of self-interest is quite common, individuals are often oriented toward cooperation, trying to maximize both their own outcomes and those of a partner. Another frequent motive is competition, where individuals, rather than trying to maximize their own gain, attempt to maximize the difference between their own outcome and the outcome of a partner. In opposition to this is an equalitarian orientation, where individuals try to keep their own outcomes and the outcomes of the partner even. Value orientation depends on personality and other individual characteristics of the decision-maker and the relation between the decision-maker and the partner, but also the situation, cultural considerations, and the nature of the decision.

2.2.2 Emotions

Human decisions and judgments are also strongly shaped by emotions. Common sense analysis of market dynamics is often made in terms of the dominance of fear or greed. Many lines of research in psychology have proven that emotions influence memory, information processing, judgments, and decisions. Positive emotions, for example, facilitate action and risky decision-making, while negative emotions prompt individuals to refrain from acting and encourage safety-seeking. The emotional congruency effect describes the tendency to recall positive memories when individuals experience positive emotions, and to recall negative memories when in a negative mood.

There is also evidence both from psychology [35, 107, 154] and neurophysiology that decisions are often made on the basis of emotions, while cognitive processes serve to justify the decisions.

2.2.3 Self-Structure

Self-structure is the largest and most powerful of psychological structures. It encodes all self-relevant information and performs regulatory functions with respect to other psychological structures. It is also a strong source of motivation. Two motives dominate the regulatory functions of self-structure. The self-enhancement motive drives the tendency for positive self-evaluation. Conforming to social norms, seeking approval, striving to win against the competition, or trying to look attractive all stem from the self-enhancement motive.

People also have a tendency to confirm their view of themselves. This tendency is described as self-verification. If someone believes that he or she cares about the

environment, he or she would likely donate to environmental causes. If someone believes he or she has no talent at math, this person is likely to do poorly in this subject. In other words, people are likely to engage in actions congruent with beliefs they have about themselves.

Beliefs concerning the self may originate in feedback from others. This process is called labeling. If others perceive an individual as a risky player, for example, the individual will be more likely to take risks. What is interesting is that people have the tendency to construct beliefs of themselves on the basis of observing their own actions. In effect, if individuals for some reason engage in a behavior, e.g., trading, this may lead to the development of a self-view as a trader, and this, in turn, may encourage them to engage in further trading.

The two modes of functioning of self-structure are called promotion and prevention [63]. In the promotion mode, individuals are oriented toward achievement. They are seeking situations where they can succeed, be noticed, win a competition, etc. Their emotions oscillate between the happiness of success and the sadness or anger associated with failure. In the prevention mode, individuals aim to minimize the losses. Everything that is happening around them is perceived as a potential threat. Individuals have the tendency to refrain from action since an action could potentially result in negative consequences. Their behavior is aimed at self-protection. Their emotions vary between anxiety and relaxation. Whether an individual is in the promotion or prevention mode depends to some extent on the individual's own general characteristics, but also on the situation and their current mood. Factors that are highly valued in promotion mode may have a negative effect in prevention mode.

Economic theory assumes that rational individuals take into account all the relevant information. Research in psychology shows that the decision-making process consists of different stages. Up until a certain moment, individuals are open to all incoming information. When they arrive at a decision, however, their information processing changes qualitatively. They selectively seek information that supports their decision, and avoid information that contradicts their decision. They are no longer open to arguments. This phenomenon is called cognitive closure.

Individuals differ in the strength of their tendency for cognitive closure [78]. Individuals characterized by a high tendency for cognitive closure form opinions relatively soon, after hearing a small amount of information. From this moment they concentrate on defending the opinion they have already formed. Individuals with a low tendency for cognitive closure do not form fixed opinions for a relatively long time. They are open to new incoming information and arguments, and are ready to adjust their opinions.

The tendency for cognitive closure also depends on the situation. Situations in which it is difficult to concentrate, for example, with a lot of noise, encourage early cognitive closure. Cognitive overload, i.e., trying to perform multiple mental operations at the same time also tends to accelerate cognitive closure. Time pressure has similar effects. In summary, although information may be available, some individuals under some circumstances are likely to ignore information that is contrary to an opinion they have already formed.

Deviations from rationality on the individual level are not necessarily incompatible with the existing theories in economy. Although each individual's decisions are to some extent irrational, it may be that these deviations from rationality can be treated as errors or noise, and while this may be quite pronounced on the individual level, the error of many individuals taken together on the group level could cancel out, provided that the individual errors are not correlated, i.e., provided that they are independent of each other.

If the errors of different individuals are correlated, as for example in the case of false information influencing a large proportion of traders, individual errors could add up, rather than cancel. In this case the predictions of the rational model may still be corrected by adjusting the outcome in a way that reflects the influence of the external factor. Social processes occurring among individuals can make the outcome of a group process very different from the sum of the outcomes of individual processes. Processes occurring at the group level cannot be reduced to processes occurring at the individual level.

2.2.4 Biases

In the following, we introduce the most commonly reported and empirically verified human biases. The list of biases given here is by no means meant to be exhaustive and we will only give a rather crude introduction to the topic.

Framing. Framing refers to the cognitive bias in which people make different decisions depending on the way a problem is presented to them. If one frames the same question in different ways, people will give different answers, even though the question the people answered was the *same* question [74]. The example in [141] gives a clear illustration of this bias. In a thought experiment, participants were offered two different solutions to save lives in a situation where 600 people had contracted a deadly disease. In the first solution labelled A, they were offered two choices:

- Save 200 people's lives out of the 600.
- Save 600 people's lives with a chance of 33 and 66 % of saving no one.

In the second solution B the participants were offered exactly the same scenario but described differently:

- Let 400 die, but save 200.
- With a 33 % chance you save all 600 lives, but with a 66 % chance that 600 people will die.

Given the way the problem is presented in solution A, people will in general opt for the first choice, because the second seems risky. However, in solution B, people will instead opt for the second choice. In this description, it is the hope of saving people that makes the second choice attractive. Notice that the two solutions A and B are the same but framed differently, through a different use of words.

Overconfidence. Extensive evidence shows that people are overconfident in their judgments. Two different but related forms appear:

- The confidence intervals that people assign to their estimates of quantities are usually far too narrow. A typical example is the case where people try to estimate the level of a given stock market index a year from today.
- People are very poor when it comes to actually estimating probabilities. We will discuss this point in more detail below, when discussing prospect theory.

Law of Small Numbers. This judgmental bias happens because individuals assume that the characteristics of a sample population can be estimated from a small number of observations of data points. According to the law of *large* numbers, the probability distribution of the mean from a large sample of independent observations of a random variable is concentrated around the mean with a variance that goes to zero as the sample size increases. In contrast, the law of *small* numbers in psychology [73] describes the bias people introduce in neglecting the fact that the variance for a small sample is larger than the variance for a big sample. One manifestation of this bias is the belief that, in a small or a large hospital, it is equally likely to find a daily birth rate of girls in excess of 50 %. Because the variance is greater for a small hospital, it is actually more likely to find a birth rate of girls in excess of 50 % at a small hospital as compared to finding the same event at a larger hospital.

Self-Attribution Bias and Hindsight Bias. Overconfidence may in part stem from two other biases, self-attribution bias and hindsight bias:

- Self-attribution refers to people's tendency to ascribe any success they have in some activity to their own talent, while blaming failure on bad luck.
- Hindsight bias refers to people's tendency to believe, after an event has occurred, that they predicted it before it actually happened.

Optimism and Wishful Thinking. Most people display an unrealistically rosy view of their abilities and prospects. Typically, over 90 % of those surveyed think they are better than the average in skills such as driving (whereas a survey in an unbiased population should give 50 %). People also display a systematic planning fallacy: they predict that tasks such as writing a book will be completed much sooner than they actually are, and that the book will be understood by the reader much more than may actually be the case!

Belief Perseverance. There is much evidence that once people have formed an opinion, they have the tendency to cling to it too tightly and for too long. In this respect two different effects appear to be relevant. The first effect is that people are reluctant to search for evidence that contradicts their own beliefs. Secondly, even if they actually do find such evidence, they treat it with excessive skepticism.

Anchoring. Anchoring is a term used in psychology to describe the common human tendency to rely too heavily on, or to anchor onto, one piece of (often irrelevant) information when making decisions. In a later section of this chapter, we will come back to the term in more detail, and give a specific recipe for identifying anchoring in financial markets.

2.3 Prospect Theory

The following description follows the text that accompanied the 2002 Nobel Prize in Economics attributed to Daniel Kahneman. We will, however, try to concentrate on the essence of the theory, rather than give a general overview. As we will see, the core idea behind prospect theory comprises three elements which are all new with respect to standard economic theory. However, before we introduce these three new elements, we first give a formal representation in the box below, in order to understand the differences. The reader who is not interested in the formal description can jump directly to the explanation given after the box.

When it comes to human decision-making, standard economic theory relies heavily on the idea that each decision-maker tries to maximize her or his utility. If we call the utility function u , then such a function is defined on a set of possible outcomes $X = \{x_1, x_2, \dots, x_N\}$. Assume for simplicity that the decision-maker has to choose between two different actions a and b . Let p_i be the probability for x_i which results in the wealth w_i under the action a , and q_i the probability for the same outcome and wealth, but instead under the action b . Then classical economic theory gives the following criterion for choosing a over b :


$$\sum_i p_i u(w_i(x_i)) > \sum_i q_i u(w_i(x_i)) . \quad (2.1)$$

The inequality (2.1) says that a rational decision-maker will assign probabilities to different random events and then choose the action which maximizes the expected value of her or his utility.

In contrast to this, prospect theory assumes three differences, which we first illustrate quantitatively and then explain below. In short, the three differences are (i) $w_i \rightarrow \Delta w_i$, (ii) $u \rightarrow v$, and (iii) $q_i \rightarrow \pi(q_i)$, so that instead of (2.1) prospect theory suggests:

$$\sum_i \pi(p_i) v(\Delta w_i(x_i)) > \sum_i \pi(q_i) v(\Delta w_i(x_i)) . \quad (2.2)$$


Fig. 2.1 Value assigned to gains and losses according to prospect theory. The figure illustrates loss aversion, since a small loss is assigned a much higher negative value than the positive value assigned to a gain of the same size. (Figure taken from [73])



Here follows the reasoning behind the changes imposed in prospect theory compared to standard economic theory:

- (i) $w_i \rightarrow \Delta w_i$. It is not the absolute level of wealth w that matters in decision-making, but rather the relative wealth with respect to a given reference level w_0 . In plain words, the stress you feel when investing \$1 million in a given project/financial market (expressed via the utility function) would be experienced very differently by a wealthy person like, say, Bill Gates. The reference point is often the decision-maker's current level of wealth which then gives the status quo for decision-making. Another example illustrating this point is in salary negotiations for people changing jobs. In that case the former salary will often serve as the reference point in the negotiations. However, the reference level could also be some aspirational level that a subject strives to obtain.
- (ii) $u \rightarrow v$. People are risk averse. This is illustrated in Fig. 2.1, which shows that the value that people assign to gains versus losses is *not* symmetric. As can be seen from this figure, even the slightest loss is something people really dislike, whereas the same amount of gain is not something considered highly desirable. This difference is often phrased by saying that the utility function has a 'kink' for small losses. Interestingly, it has recently been pointed out that financial markets themselves react in such a risk averse manner, with the correlations between stocks in an ascending market being different from the correlations between stocks in a descending market. For a discussion on this point, see, e.g., [13, 114].
- (iii) $q_i \rightarrow \pi(q_i)$. The last part of prospect theory relates to the fact that people have problems assigning the proper probability to events, often placing too high a weight on events that are highly unlikely to occur and placing too little weight on events that are very likely to occur. This is illustrated in Fig. 2.2. We only have to think about lottery tickets: people keep on buying them even though their chance of winning the jackpot is basically nil.

Fig. 2.2 Decision weight as a function of probability. The figure illustrates the fact that people assign too high a weight to events that are very unlikely to occur and too little weight to events that are almost sure to occur. (Figure taken from [73])



In our view what is particularly appealing about prospect theory is that it can actually be checked, and it can make predictions about how people will behave in different situations, which can then be used to cross-check the theory. In other words, this is an example of a falsifiable theory, something we will return to and explain in Chap. 4. Let us just finish this section by mentioning that the three assumptions of prospect theory have indeed been checked in many experiments and under various decision-making conditions.

2.4 Pricing Stocks with Yardsticks and Sentiments

Having discussed the way certain biases can influence human decision-making, it is time to suggest in practical terms how biases could weight the way stocks find their prices in the market. The aim in this section will thus be to ‘quantify’ human sentiments and show how, in certain cases, they can influence the pricing of a given stock. We will suggest the somewhat surprising result that the pricing of a given stock can be expressed in terms of the general ‘sentiment’ of the market. This is a very similar situation to one we discussed in Sect. 1.5, except that in that case the pricing of a given stock was expressed in terms of, not the sentiment, but the general performance of the market. It turns out that the formula we find for the pricing of a given stock in terms of sentiment has a very similar structure to that of the CAPM model found in Sect. 1.5.

Human decision-making by professionals trading daily in the stock market can be a daunting task. It includes decisions about whether to keep on investing or to exit a market subject to huge price swings, and how to price in news or rumors attributed to a specific stock. The question then arises as to how professional traders, who specialize in the daily buying and selling of large amounts of a given stock, know how to price a given stock properly on a given day? Here we introduce the idea that people use heuristics, or rules of thumb, with reference to certain ‘yardsticks’ derived from the performance of the other stocks in a stock

index. Under- or over-performance with respect to such a yardstick then signifies a generally negative or positive sentiment of market participants towards a given stock. Using the empirical data from the Dow Jones Industrial Average, stocks can be shown to have daily performances with a clear tendency to cluster around the measures introduced by these yardsticks. We illustrate how sentiments, most likely due to insider information, can influence the performance of a given stock over a period of months, and in one case years.

2.4.1 Introduction

One of the founders of behavioral finance, D. Kahneman (Shefrin [123]), once pointed out how media coverage of financial markets tends to depict them with the traits of a stereotypical individual. Indeed, the media often describe financial markets with attributes like “thoughts, beliefs, moods and sometimes stormy emotions. The main characteristic of the market is extreme nervousness. It is full of hope one moment and full of anxiety the next day.” One way to get a first quantification of the sentiment of the market is to probe the sentiments of its investors. Studying sentiments of consumers/investors and their impact on markets has become an increasingly important topic [146]. Several sentiment indices of investors/consumers already exist, and some have now been recorded over a time span of a few decades. The Michigan Consumer Sentiment index, published monthly by the University of Michigan and Thomson Reuters, is probably the one which has the largest direct impact on markets when published. The natural question then arises as to whether it is possible to predict market movements by considering the sentiments of consumers/investors?

Fisher and Statman [45] made a study of tactical asset allocation from data describing the sentiment of a heterogeneous group (large, medium, small) of investors. The main idea in [45] was to look for indicators for future stock returns based on the diversity of sentiments. The study found that the sentiments of different groups do not move in lockstep, and that sentiments for the groups of large and small investors could be used as contrary indicators for future S&P 500 returns. However, more recent research [135] on investor sentiment expressed in the media (as measured from the daily content of a *Wall Street Journal* column) seems to point in the opposite direction, with high media pessimism predicting downward pressure on market prices. Such results are more in line with theoretical models of noise and liquidity traders [38, 39]. Other studies [8] claim very little predictability of stock returns using computational linguistics to extract sentiments on 1.5 million Internet message boards posted on Yahoo! Finance and Raging Bull. However, in that study it was shown that disagreement induces trading and also that message posting activity correlates with volatility of the market.

Common to all the aforementioned studies is the aim to predict global market movements from sentiments obtained either from surveys or from internet message boards. In the following, we propose instead to obtain a sentiment-related pricing for a given asset by expressing the sentiment of a given stock *relative* to the market.

This is similar to the principle of the Capital Asset Pricing Model (CAPM) which we met in Sect. 1.5, since it relates the price of a given asset to the price of the market, instead of trying to give the absolute price level of the asset/market. Put differently, we will in the following introduce a method that does not estimate the impact that a given ‘absolute’ level of sentiment can have on the market, but instead introduce a sentiment of an asset *relative* to the sentiment of the general market, whatever the absolute (positive/negative) sentiment of the general market. As we shall illustrate, this gives rise to a pricing formula for a given stock relative to the general market, much like the CAPM, but now with the relative sentiment of the stock to the market included in the pricing.

2.4.2 Theory of Pricing Stocks by Yardsticks and Sentiments

In the following we consider how traders find the appropriate price level of a given stock on a daily basis. One could for example have in mind traders that specialize in a given stock and actively follow its price movements so as to determine opportune moments either to buy certain amounts or instead to sell as part of a larger order. The question is, what influences the decision-making for traders as to when to enter and when to exit the market? As we saw in Chap. 1, according to the standard economic view, only expectations about future earning/dividends and future interest rate levels should matter in the pricing of a given stock. Looking at the often very big price swings during earnings or interest rate announcements, this part clearly seems to play a major role, at least in some specific instances. But what about other times when there is no news which can be said to be relevant for future earnings/interest rates? The fluctuations seen in daily stock prices simply cannot be explained by new information related to these two factors, nor can risk aversion, so why do stock prices fluctuate so much, and how do traders navigate the often rough seas of fluctuations?

Here we take a heuristic point of view and argue that traders need some rules of thumb, or as we prefer to say, yardsticks, in order to know how to position themselves. In the box below we will derive a new pricing formula for stocks based on the relative ‘performance’ of a stock with respect to the ‘performance’ of the market. The performance we suggest is basically the ratio of the return of a stock to its risk, where we quantify risk in terms of how volatile a stock is. The main idea is that professionals, possibly with insider information, may have a bias with respect to the stock they trade, and this bias will make their trade weigh on the performance of a stock relative to the general market.

A first rough estimate for a trader would obviously be the returns of other stocks in a given stock index. Let $s_i^I(t)$ be the daily return of stock i belonging to index I at time t , and $R_{-i}^I(t)$ the return of the remaining $N - 1$ stocks in the

(continued)

(continued)

index I at time t . We emphasize the *exclusion* of the contribution of stock i in R_{-i}^I in order to avoid any self-impact which would amount to assuming that the price of a stock rises because it rises. We use a capital letter to denote the index and a lower case letter to denote a specific stock i . Using the average return of the other stocks as a first crude yardstick, one would then conclude that traders of stock i would price the stock according to

$$s_i^I(t) \approx R_{-i}^I(t) \equiv \frac{1}{N-1} \sum_{j \neq i} s_j^I(t). \quad (2.3)$$

A powerful tool that is often used in physics to check the validity of an equation is dimensional analysis, checking that the quantities on each side of the equation have the same dimensions. We will give a more detailed explanation of the method in Chap. 5. By the same token, an expression should be independent of the units used. Since (2.3) expresses a relationship between returns, i.e., a quantity that expresses increments in percentages, it is already dimensionless. However, we argue that there is a mentally relevant ‘unit’ in play, namely the size of a typical daily fluctuation of a given stock. Such a mental ‘unit of fluctuation’ is created by the memory of traders who closely follow the past performance of a given stock. Dividing both sides of (2.3) by the size of a typical fluctuation would therefore be one way to ensure independence from such units. Taking the standard deviation as measure, the renormalized (2.3) takes the form

$$\frac{s_i^I(t)}{\sqrt{\langle \sigma^2(s_i^I) \rangle_T}} = \frac{R_{-i}^I(t)}{\sqrt{\langle \sigma^2(R_{-i}^I) \rangle_T}}, \quad (2.4)$$

where the variable $\sigma^2 \equiv \langle X^2 \rangle - \langle X \rangle^2$ denotes the variance of the variable X and $\langle \rangle_T$ denotes an average over a given window size T .

As we will show in a moment, (2.4) turns out to be a good approximation for most stocks over daily time periods. There are, however, strong and persistent deviations. Actually, in the following we will define stocks for which (2.4) holds on average as *neutral* with respect to the sentiment traders have on the given stock. Similarly, we use this relation as a measure, positive or negative, of how biased a sentiment traders have on the given stock. More precisely, the sentiment α_i^I of a given stock i is defined as

$$\alpha_i^I(t) = \frac{s_i^I(t)}{\sqrt{\langle \sigma^2(s_i^I) \rangle_T}} - \frac{R_{-i}^I(t)}{\sqrt{\langle \sigma^2(R_{-i}^I) \rangle_T}}. \quad (2.5)$$

(continued)

(continued)

We emphasize that the sentiment is defined with respect to the other stocks in the index, which serve as the neutral reference. The ratio of a stock's (excess) return to its standard deviation tells us something about its performance, or reward-to-variability ratio, also called the Sharpe ratio in finance. Therefore, (2.5) attributes a positive (resp. negative) bias/sentiment to a stock, $\alpha_i^I > 0$ (resp. $\alpha_i^I < 0$), when the Sharpe ratio of the stock exceeds (resp. underperforms) the Sharpe ratio of the sum of the other stocks in the index [134].

Rewriting (2.5), the pricing of stock i can now be given in terms of a renormalized performance of the other stocks in the index as well as a possible bias:

$$s_i^I(t) = \sqrt{\langle \sigma^2(s_i^I) \rangle_T} \alpha_i^I(t) + \frac{\sqrt{\langle \sigma^2(s_i^I) \rangle_T}}{\sqrt{\langle \sigma^2(R_{-i}^I) \rangle_T}} R_{-i}^I(t). \quad (2.6)$$

As a first check of (2.6), we take the expectation value $E(\cdot)$ of (2.6) by averaging over all stocks listed on the index, and then average over time (daily returns). A priori, over long periods of time, one would expect to find as many positively biased as negatively biased stocks in an index composed of many stocks. Using this assumption the term in α_i^I disappears by symmetry. One obtains

$$E(s_i^I) = \frac{E(R_{-i}^I)}{\sqrt{\langle \sigma^2(R_{-i}^I) \rangle_T}} \sqrt{\langle \sigma^2(s_i^I) \rangle_T}. \quad (2.7)$$

Equation (2.7) is very similar in structure to the Capital Asset Pricing Model in Finance (CAPM) [150], discussed in Sect. 1.5:

$$\frac{E(s_i^I) - R_f}{\beta_i} = E(R^I) - R_f, \quad \beta_i = \frac{\text{Cov}(s_i^I, R^I)}{\sigma^2(R^I)}, \quad (2.8)$$

where R_f in (2.8) is the risk-free return which, since we consider daily returns, will be taken equal to 0 in the following:

$$E(s_i^I) = \frac{\text{Cov}(s_i^I, R^I)}{\sigma^2(R^I)} E(R^I). \quad (2.9)$$

The main difference between the CAPM in the form (2.9) and our hypothesis (2.7) is that we stress the use of standard deviations in the pricing formula, rather than the covariance between the stock return and the index return on the right-hand side of (2.9). Furthermore, we argue that the covariance between a given stock




Fig. 2.3 Testing different hypotheses. Data showing pricing according to the CAPM hypothesis (*left*) and the sentiment hypothesis (*right*). If the data follows one of these hypotheses, one should see it spread evenly around the diagonal (*fat solid line*) in the corresponding plot. This seems to be the case for the sentiment pricing, but not for the CAPM pricing, where data points are ‘tilted’ with respect to the diagonal. The plot *on the left* shows the CAPM hypothesis (2.9) using the daily returns of the Dow Jones Industrial index over the period 3 January 2000 to 20 June 2008. The plot *on the right* illustrates our hypothesis (2.7) using the same data set. Each point corresponds to a daily return s_i of a given stock i

and the index is not a very stable measure over time, in contrast to the variance of a given stock. One cause of instability in the covariance could for example be sudden ‘shocks’ in terms of specific good or bad news for a given company. After such a shock, we postulate that the covariance between the stock and the index changes, whereas the stock’s variance remains the same, but with a change in relative performance. The pricing formula that we obtain is reminiscent of the so-called capital allocation line in finance. This expresses the return of a portfolio that is composed of a certain percentage of the market portfolio, but with the remainder invested in a risk-free asset. However, the capital allocation line only expresses the return of this specific portfolio, whereas our expression is supposed to hold true for each individual asset.

The data points in Fig. 2.3 are used to contrast the CAPM hypothesis and our hypothesis using daily returns of the Dow Jones Industrial Average over almost a decade of data [136]. A perfect fit of the data to the two equations would in each case lie on the green diagonal. The data for CAPM appear tilted with respect to the diagonal, whereas the data concerning our hypothesis appear to be symmetrically distributed around the diagonal, which is what one would expect if the data




Fig. 2.4 Extracting sentiment biases in stocks. *Left:* Cumulative returns of four different stocks in *blue* over a given period (see below) compared with the return of the general market index (Dow) in *green* over the same period. *Right:* Corresponding cumulative bias in sentiment for the given stock, plotted in *red*. The fact that a constant decline in the cumulative sentiment can be observed over a certain period (for Citibank stock over the whole period) indicates the underperformance of the stock with respect to the general market. The data for Citibank, Caterpillar, and United Technologies Corporation are for the time period 3 January 2000 to 20 June 2008, and the data for Cisco are for the time period 1 January 2009 to 2 June 2011

obeyed (2.7) on average. The fact that the cloud of data points is symmetrically scattered around the diagonal gives the first evidence in support of (2.7).

The sentiment α in (2.5) was introduced as a behavioral trait, and as such we would expect to see its effect on a long time scale of at least the order of weeks or months. Figure 2.4 shows the cumulative sentiment for four different stocks, Citibank, Caterpillar, and United Technologies Corporation over the period from 1 June 2000 to 20 June 2008, and Cisco over the time period from 1 June 2009 to 2 June 2011. The plots to the left show in green the return of the Dow Jones and in blue the given stock over the given time period.

The case of the Citibank stock is particularly striking, with a constant negative sentiment seen by the continuous decline in the cumulative sentiment curve of Fig. 2.4, corresponding to a constant underperformance over 2 years. It should be noted that the data was chosen in order to have both a declining general market, as happens over the first half of the period shown, and also an increasing overall market, as happens over the rest of the time period chosen. It is remarkable that the sentiment of the Citibank stocks remains constant regardless of whether the general trend is bullish or bearish.




Fig. 2.5 Probability distribution function of the sentiment variable α_i

Similarly, it should be noted that the general sentiment for Caterpillar and United Technologies Corporation was neutral in the declining market, but then became distinctly negative over the last 3 or 4 months of the time series when the general market was bullish. The price history for Cisco Systems tells a similar story. The only difference here is the two big jumps occurring after 350 and 400 days. These two particular events took place on 11 November 2010 and on 10 February 2011. On 11 November 2010, the price dropped because of a bad report for the third quarter earnings. This gave rise to a loss of confidence by investors who were expecting a sign of recovery after a couple of hard months. On 10 February 2011, Cisco Systems announced a drop in their earnings (down 18 %) together with a downward revision (−7 %) of sales of their core product.

It is worth noting that the decline in cumulative sentiment took place *before* the two events: prior to 11 November 2010, there was a long slow descent of the cumulative sentiment (implying a constant negative sentiment), and after the 11 November 2010 the descent continued. This could be taken as evidence that some investors with insider knowledge were aware of the problems of the company, which was revealed only to the public on the two aforementioned days.

Figure 2.5 shows the probability distribution function of the sentiment variable α defined in (2.5), obtained by sampling, using the daily return for all stocks of the Dow Jones Industrial Average over the period 3 January 2000 to 20 June 2008. As can be seen from the inset of Fig. 2.5, the distribution appears to be exponential for both positive and negative sentiments. One notes that the empirical distribution is

symmetric with respect to the sign of the sentiment, something which was implicitly assumed in deriving (2.7) from (2.6).

2.4.3 Discussion

To sum up, we have pointed out the importance of a measure of the sentiment of a given stock *relative* to its peers. The idea is that people use heuristics, which one might call rules of thumb or yardsticks, obtained from the performance of the other stocks in a stock index. Under- or over-performance with respect to a yardstick then signifies a general negative or positive sentiment of the market participants towards a given stock. We have introduced a quantitative method to check such an assumption.

The bias created in such cases does not necessarily have a psychological origin, but could be due to insider information. Insiders having superior information about the state of a company reduce/increase their stock holding, gradually causing a persistent bias over time. The introduction of a measure for the relative sentiment of a stock has allowed us to come up with a pricing formula for stocks very similar in structure to the CAPM model. Using empirical data from the Dow Jones Industrial Average, stocks are shown to have daily performances with a clear tendency to cluster around the measures introduced by the yardsticks, in accordance with our pricing formula.

2.5 Sticky Price Dynamics: Anchoring and Other Irrational Beliefs Used in Decision Making

The last section gave an example of how to quantify price formation in financial markets due to a psychological phenomenon caused by people using heuristics to price a stock relative to an index. We will continue along similar lines in this section and look at how irrational beliefs can influence the pricing of stocks. We will then introduce a tool to detect and quantify this. The aim is to take seriously some of the ideas coming from behavioral finance, but also to introduce quantitative methods which can provide us with tests. The tests should be directly applicable to financial market data in order to verify the presence or otherwise of such irrational human beliefs in the markets.

Even though we limit ourselves in this chapter to the impact of psychological factors, the approach should be seen as more general, suggesting a way to introduce quantitative methods from psychology and sociology that will shed light on the impacts of such factors on the pricing of assets. The ideas coming from behavioral finance should not therefore be seen as limiting, but rather as a starting point for tackling the more general question of the impacts of sociological and psychological factors on the markets.

We end this chapter by giving yet another example of how to extract information at the individual level of a representative agent trading in the market, before




Fig. 2.6 How to make particles move in a flashing ratchet and how to use this idea in a trading algorithm. A flashing ratchet switches between an on state and an off state. In the on state, Brownian particles concentrate around the potential minima. In the off state, the particles undergo one-dimensional isotropic diffusion. Periodic or stochastic switching between the two states induces a particle current in the positive x direction. As explained in the text, a patient trader can use ‘sticky’ price movements and wait for a favorable moment to enter the market, just as the particles in the figure have to wait for the ratchet to be on to make a ‘favorable’ or positive move (The figure is taken from [85])

discussing ways to study effects at the group (sociological) level in the next chapter. It is not immediately obvious that the way motors work in biology could be relevant to financial trading. In certain situations, this will nevertheless be our claim. More specifically, we will use insights into the way biological motors work, to suggest ways to detect the presence in financial markets of a specific human behavioral trait known as anchoring.

To understand the idea, consider Fig. 2.6, which shows an example of a so-called flashing ratchet. The flashing ratchet has been suggested as a mechanism that provides some biological molecules with motors allowing them to move and operate on a micro- to nanoscopic scale. The ratchet flashes, say randomly, between on and off, and in doing so, this induces a drift of the red particles, as illustrated in the figure. The basic mechanism behind this is brought about by the asymmetric potential of the ratchet, which pulls the particles to the right whenever the ratchet is on. When the ratchet is off, the particles diffuse freely without any average drift. Another way of looking at the same problem is to say that the combined system of the ratchet and particles acts as if particles were influenced by a motor. Looking at it this way, a particle just waits, and whenever there is an opportune fluctuation (the ratchet is on), it exploits this fluctuation to move.

The reason for discussing this mechanism is the a priori surprising suggestion of using this same mechanism in a trading algorithm. This idea, proposed by Lionel Gil [57], gives yet another startling example of how cross-disciplinary ideas can bring unexpected new insights into a problem where no one would have expected a link. We will describe the idea in detail shortly, but first we need to discuss the behavioral trait known as anchoring, introduced by Kahneman and Tversky in the 1970s.

Anchoring is a term used in psychology to describe the common human tendency of relying too heavily on (being anchored to) one piece of often irrelevant information when making decisions. One of the first observations of anchoring was reported in a now classic experiment by Tversky and Kahneman [140]. Two test groups were shown to give different mean estimates of the percentage of African nations in the United Nations depending on a completely unrelated (and randomly generated) number suggested by the experimenters to the two groups. Evidence for the human tendency to anchor has since been reported in many completely different domains. For example, customers show inertia when it comes to switching from a given brand [152]. In this case it is the old brand's price that acts as an anchor. Other evidence comes from studies of online auctions [41]. People have a tendency to bid more for an item the higher the 'buy-now' price. Anchoring has also been shown in connection with property prices [102]. In this case it was shown that a subject's appraisal depends on an arbitrarily posted house price.

In the context of financial markets, anchoring has been observed via the so-called disposition effect [62, 75, 104, 124], which is the tendency to sell assets that have gained value and keep assets that have lost value. In that case the *buying price* acts as an anchor. This is different from the anchoring discussed in the following, where a recent *price level* acts as an anchor. As noted in [148], conclusive tests using real market data are usually difficult because investors' expectations, as well as individual decisions, cannot be controlled or easily observed. However, in experimental security trading, subjects have been observed to sell winners and keep losers [148].

The main premise in the following is that situations occur in financial markets that can lead to 'sticky' price dynamics, and that this can be detected via the flashing ratchet method. One can think of a variety of circumstances that could create such 'sticky' price dynamics. The European Monetary System (EMS) gives a particularly clear example. In this case a monetary policy induces bands over which currencies were allowed to fluctuate with respect to one and another, illustrating a policy-caused mechanism for such 'sticky' price movements in financial markets [57].

As we shall see shortly, another possible situation leading to 'sticky' price movements in equity markets is when market participants actively follow the price over days or weeks, thereby creating a subjective reference (anchor) and memory of when an asset is 'cheap' or 'expensive'. Several studies on the persistence of human memory have reported sleep as well as post-training wakefulness before sleep to play an important role in the offline processing and consolidation of memory [109]. It therefore makes sense to think that conscious as well as unconscious mental processes influence the judgment of those who specialize in active trading on a day-to-day basis.

The idea behind the flashing ratchet method can now be outlined as follows. We assume that anchoring is present in financial markets because of 'sticky' price movements, with recent prices used by market participants as an anchor in determining whether an asset is over- or undervalued. Such irrational behavior would in principle open up arbitrage possibilities for speculators buying when an asset is perceived to be undervalued and selling when it is overvalued. As will be

Table 2.1 Table showing the four different configurations x_1, x_2, x_3, x_4 that can occur when trading two assets in the ‘sticky price’ algorithm. Quasi-static price levels of the two assets are denoted by \bar{A}_1 and \bar{A}_2 and the fluctuations around those levels are dA_1 and dA_2 , respectively

x_1	x_2	x_3	x_4
A_1 —	A_1 —		
$\Downarrow dA_1$	$\Downarrow dA_1$		
\bar{A}_1 —	\bar{A}_1 —	\bar{A}_1 —	\bar{A}_1 —
		$\Downarrow dA_1$	$\Downarrow dA_1$
		A_1 —	A_1 —
	A_2 —		A_2 —
	$\Downarrow dA_2$		$\Downarrow dA_2$
\bar{A}_2 —	\bar{A}_2 —	\bar{A}_2 —	\bar{A}_2 —
$\Downarrow dA_2$		$\Downarrow dA_2$	
A_2 —		A_2 —	

seen, just as the biological motors make a move when a favorable fluctuation occurs, so a patient trader can wait and act when the right price fluctuation occurs.

Table 2.1 illustrates how the flashing ratchet algorithm works in the simplest case with only $N = 2$ different assets. For the interested reader, a complete analytical derivation of the problem is given for this simple case in the appendix in Sect. 2.5.1. The table illustrates the different configurations that can arise with two assets, assuming price fluctuations around quasi-static price levels given by \bar{A}_1 for asset 1 and \bar{A}_2 for asset 2. The solution of the method can easily be generalized to an arbitrary number of assets and with price levels which, instead of remaining constant, vary slowly over time. To simplify things, we will only consider the case where a trader always takes the long position for one asset. The general case for short only or short plus long positions is easy to figure out. It will be useful to keep Table 2.1 in mind in the following.

Here is how the flashing ratchet idea works. Assume that you are long on the position of say asset 1 bought at an earlier time. Let us also assume for simplicity that the share of asset 1 was bought at the reference price level given by \bar{A}_1 . As illustrated in Table 2.1, four different situations can occur, corresponding to the four different configurations x_1, x_2, x_3, x_4 , but your reaction to the four different cases should be different, as we shall illustrate. Consider first that, e.g., configuration x_3 occurs, that is, we have a fluctuation where the price A_1 of asset 1 is below the quasi-static price level \bar{A}_1 by an amount dA_1 , and that simultaneously the situation is similar for asset 2. This is not a profitable situation for you, so if configuration x_3 occurs and you are long on asset 1, you accept your temporary loss and do nothing. A patient trader would instead wait till configuration x_1 occurs, since in this case asset 1 is overvalued and asset 2 is undervalued, giving the opportunity to close the long position on asset 1 and open a new long position on asset 2. Now being long on asset 2, a patient trader would have to wait until configuration x_4 comes up, closing

the long position on asset 2 and opening a new long position on asset 1, and so on and so forth.

In principle, this sounds like a very simple algorithm for statistical arbitrage, but in order to be able to settle the question as to whether such a strategy is profitable, and evaluate how risky it would be, one would need to know three things:

- The probability of finding the price of the asset to have a given value A . The method assumes the price evolution to be ‘sticky’, i.e., quasi-stationary, so that at any given time t a market participant has a notion of the likelihood of a certain price value A at that very moment of time. It would be natural for this probability to change over time, but to keep the discussion simple, it will be assumed in the following to be time-independent. Generalization to a time-dependent probability distribution is straightforward. In short one needs to know the price probability distribution function $P(A)$.
- The probability $P(A \rightarrow B)$ of going from one quasi-stationary price A to another quasi-stationary price B .
- The transaction costs C .

The method is described in detail in Sect. 2.5.1 and quantifies the circumstances under which the presence of anchoring in financial markets would be detectable in terms of a profitable investment strategy that speculates on this particular cognitive bias. The method will first be illustrated in the case where an exact solution is available, before applying the algorithm to real financial data.

To see how the method works in a controlled setup, we assume a fictitious market with quasi-stationary price levels \bar{A}_i , and for simplicity take $\bar{A}_1 = \bar{A}_2 = 1.0$. Assume further that the price fluctuations have only two possible values $dA_1 = dA_2 = \pm 0.11$ with equal probability for the plus or minus signs. Note that Table 2.1 gives a very simple schematic representation of the problem, which is by its very nature probabilistic. In general, anyone observing the market will see a fluctuating time series of prices with no a priori knowledge of the underlying probability distributions governing the price dynamics. As mentioned before, one can also imagine markets in which price probability distributions change slowly over time. In order to use the flashing ratchet algorithm, a market observer then has to estimate the quasi-static levels A_i . There are various ways to do so. One of the simplest is just to take the average of the prices over the last m days.

The solid line in Fig. 2.7 shows the steady state analytical solution given by the average return of the algorithm. The figure shows the performance of the algorithm (circles) as a function of time. After some initial fluctuations in performance, the algorithm is indeed seen to converge to the theoretical average return value. The general performance of the algorithm in terms of the expected average return and the expected average risk (taken here as the standard deviation of the return) can be found by using (2.14)–(2.16), given in the Sect. 2.5.1. To see how the numerical flashing ratchet algorithm checks against these expressions, the average return and average volatility are shown in Fig. 2.8 against one of the model parameters, viz., dA_1 , for the simple case of a Bernoulli distribution (2.18) and (2.19).

It should be noted that the results of [57] and those presented here derive from an effect also mentioned in [130], which appears as soon as the price exhibits mean




Fig. 2.7 Return of trading algorithm. Averaged return $r(t)$ as a function of time t in the trading algorithm that exploits sticky price dynamics. We test the algorithm against the simplest possible case where the price fluctuates by $\pm dA$ around a given price level \bar{A} . Circles represent the result obtained by using the algorithm (2.16) with $\bar{A}_1 = \bar{A}_2 = 1.0$ and $dA_1 = dA_2 = 0.11$, and memory $m = 5$. The solid line represents the analytical expression (2.18)

reversal. However, in [130], just one asset was considered, and no calculation was done with respect to risk and return for a portfolio. Apart from the policy-induced case of the EMS, there is no a priori reason to expect real market data to show truly ‘quasi-static’ behavior on longer time scales of months or years. Problems relating to such long term drifts in the price anchor \bar{A}_i were noted in [57]. When \bar{A}_i is time-dependent, the assumptions made in Sect. 2.5.1 are no longer valid, and the return of the algorithm was then shown to vanish [57].

In order to get round this difficulty, one solution might be to generalize the formalism to time-dependent probability distributions. Here, however, we suggest another approach. In principle it would be difficult to correct for the drift of a single asset or a small set of assets. But using instead a market index which is composed of a portfolio of assets, it is then possible to modify the algorithm in such a way that it is always market neutral, *regardless* of any drift in the portfolio of N assets.

Figure 2.9 gives an example of how this might work out. The figure shows the market neutral algorithm applied to real market data of the Dow Jones stock index, as well as the CAC40 stock index. The first half of the time period for the Dow Jones index was used in-sample to determine the best choice among three values of the parameter $m = 5, 10, 15$ days. The algorithm was then applied out-of-sample to the second half of the time period for the Dow Jones index, and over the whole time period for the CAC40 stock index. Since we are interested in looking at any possible




Fig. 2.8 Return and volatility of trading algorithm. Averaged return r (circles) and averaged volatility σ (crosses) versus price fluctuations dA_1 . The data points were obtained in the steady state using the flashing ratchet algorithm (2.16) with memory $m = 5$ to estimate the quasi-static price levels \bar{A}_i from which the classification was made according to Table 2.1. The random price time series (with randomness stemming from the sign of dA_i) were generated with fixed values $\bar{A}_1 = \bar{A}_2 = 1$ and $dA_2 = 0.11$. Solid lines represent the analytical results (2.18) and (2.19)

impacts coming from human anchoring, it seems justified a priori to use only daily or weekly data since the higher the frequency of trading (say seconds or minutes), the more computer dominated the trading becomes. It therefore seems reasonable to probe only three possible values corresponding to 1, 2, or 3 weeks in-sample. In order to look for arbitrage possibilities, weekly data was used in order to avoid the impact of transaction costs due to over-frequent trading.

Another reason for restricting to weekly data relates to the main claim put forward according to which market participants, by actively following the price, thereby create a subjective reference (anchor) and memory of when an asset is 'cheap' or 'expensive'. The out-of-sample profit from the market-neutral trading algorithm (with transaction costs taken into account) on the CAC40 index, as well as the second period performance on the Dow Jones index, gives evidence that anchoring does indeed sometimes play a non-negligible role on the weekly price fixing of the Dow Jones and CAC40 stock markets, and reconfirms the claim in [57] that the policy imposed by the European Monetary System leads to arbitrage possibilities. The results also give yet another illustration of the difficulty in proving market efficiency by only considering lower order correlations in past price time series.




Fig. 2.9 Trading algorithm applied to market data. Total cumulative return of the portfolio for the market neutral sticky price algorithm applied to daily price data of the Dow Jones stock index (*dotted line*) and the CAC40 stock index (*solid line*) over the period 3 January 2000 to 2 May 2006. The first half of the time period for the Dow Jones index was used in-sample to determine the best choice among three values of the parameter $m = 5, 10, 15$ days. The second half of the time period for the Dow Jones index as well as the full period for the CAC40 index were done out-of-sample with $m = 10$. A trading cost of 0.1 % was included for each transaction

2.5.1 Appendix: Quantitative Description of the Trading Algorithm

Assume that, at every time step t , an agent uses a fixed amount of his wealth to hold a long position on one out of N assets. For simplicity, $N = 2$ will be assumed in the following, but the arguments can be extended to arbitrary N . Assume furthermore that the probability distribution functions (PDFs) of the *price* of the two assets, viz., $P_1(A_1)$ and $P_2(A_2)$, are stationary distributions. Instead of the usual assumption of a random walk for the returns, short term anchoring of prices at quasi-static price levels is imposed. No specific shape is assumed and the assets may or may not be correlated, but any correlation is irrelevant for the following arguments. As noted in [57], the assumption of short term stationary prices can arise due to price reversal dynamics caused, e.g., by monetary policies. As will be argued and tested in the following, short

(continued)

(continued)

term ‘stationary’ prices can also be created by short term human memory of when an asset was ‘cheap’ or ‘expensive’.

Consider any instantaneous fluctuation of the prices (A_1, A_2) around their quasi-static price levels (\bar{A}_1, \bar{A}_2) . Classifying the 2^N different cases according to whether $A_i < \bar{A}_i$ or $A_i > \bar{A}_i$, one has (for $N = 2$) the four different configurations x_i shown in Table 2.1. In the steady state, the probability flux into a given configuration x_i equals the probability flux out of that configuration:

$$\sum_j P(x_j)P(x_j \rightarrow x_i) = \sum_j P(x_i)P(x_i \rightarrow x_j) . \quad (2.10)$$

In the steady state, the average return R_{av} per unit time is then given by

$$R_{av} = \sum_{i=1}^4 \sum_{j=1}^4 P(x_i)P(x_i \rightarrow x_j)r_{av}(x_i \rightarrow x_j) , \quad (2.11)$$

where $r_{av}(x_i \rightarrow x_j)$ is the average return gained/lost in the transition $x_i \rightarrow x_j$. For each configuration x_i , one is assumed to hold a long position of *either* asset 1 *or* asset 2. Let $s = i$ be a state variable indicating that one is long on the position of asset i . Then

$$r_{av}(x_i \rightarrow x_j) = P(s = 1|x_i)r_{av}(x_i \rightarrow x_j|s = 1) + P(s = 2|x_i)r_{av}(x_i \rightarrow x_j|s = 2) , \quad (2.12)$$

where $P(s = i|x_j)$ denotes the probability of holding asset i given that one is in configuration x_j , while $r_{av}(x_i \rightarrow x_j|s = k)$ denotes the average return in the steady state from holding asset k when there is a transition from configuration x_i to x_j , given by

$$r_{av}(x_j \rightarrow x_k|s = k) = \int dA_k \int dA'_k P(A'_k|x_i)P(A_k|x_j) \ln \frac{A'_k}{A_k} , \quad (2.13)$$

with $P(A_k|x_i)$ the *conditional* probability of getting the price A_k given that one is in configuration x_i . For example, knowing that one is in configuration x_1 , one has

$$P(A_2|x_1) = \frac{P(A_2)\theta(A_2 \leq 0)}{\int_{-\infty}^0 P(A'_2)dA'_2} ,$$

(continued)

(continued)

where $\theta(A_2 \leq 0)$ is the Heaviside function.

Using (2.11)–(2.13), the general expression for the average return gained by the algorithm takes the form

$$R_{av} = \sum_{i=1}^4 \sum_{j=1}^4 \sum_{s=1}^2 P(x_i) P(x_i \rightarrow x_j) P(s|x_i) \int dA_k \int dA_l \times P(A_l|x_i) P(A_k|x_j) \ln \frac{A_l}{A_k} . \quad (2.14)$$

The corresponding risk measured by the average standard deviation of the return is given by

$$\begin{aligned} \sigma^2 &= \langle (r - R_{av})^2 \rangle \\ &= \sum_{i=1}^4 \sum_{j=1}^4 \sum_{s=1}^2 P(x_i) P(x_i \rightarrow x_j) P(s|x_i) \int dA_k \int dA_l P(A_l|x_i) \\ &\quad \times P(A_k|x_j) \left(\ln \frac{A_l}{A_k} \right)^2 - R_{av}^2 . \end{aligned} \quad (2.15)$$

The ‘trick’ with this algorithm consists in breaking the symmetry by always choosing $P(s|x_i)$ according to the following rules:

$$\begin{aligned} P(s=1|x_1) &= 0, \quad p(s=2|x_1) = 1, \quad p(s=1|x_2) = p(s=2|x_2) = 1/2, \\ P(s=1|x_3) &= p(s=2|x_3) = 1/2, \quad p(s=1|x_4) = 1, \quad p(s=2|x_4) = 0. \end{aligned} \quad (2.16)$$

That is, if not already long, one always takes a long position on asset 2 (resp. 1) whenever configuration x_1 (resp. x_4) occurs, since the asset is undervalued in this case. Likewise, if one is long on asset 1 (resp. 2) whenever configuration x_1 (resp. x_4) occurs, one sells that asset, since it is overvalued.

To illustrate the algorithm, consider the simplest case where $P(A_i)$ takes only two values $\bar{A}_i \pm dA_i$ with equal probability 1/2. Inserting

$$\begin{aligned} P(A_1|x_1) &= \delta(\bar{A}_1 + dA_1), & P(A_1|x_2) &= \delta(\bar{A}_1 + dA_1), \\ P(A_1|x_3) &= \delta(\bar{A}_1 - dA_1), & P(A_1|x_4) &= \delta(\bar{A}_1 - dA_1), \\ P(A_2|x_1) &= \delta(\bar{A}_2 - dA_2), & P(A_2|x_2) &= \delta(\bar{A}_2 + dA_2), \\ P(A_2|x_3) &= \delta(\bar{A}_2 - dA_2), & P(A_2|x_4) &= \delta(\bar{A}_2 + dA_2), \end{aligned} \quad (2.17)$$

(continued)

(continued)

and $P(x_i) = P(x_i \rightarrow x_j) = 1/4$ into (2.14), one gets the average return:

$$R_{av}^{\bar{A}_i \pm dA_i} = \frac{1}{8} \left[\ln \frac{\bar{A}_1 + dA_1}{\bar{A}_1 - dA_1} + \ln \frac{\bar{A}_2 + dA_2}{\bar{A}_2 - dA_2} \right], \quad (2.18)$$

with variance

$$\begin{aligned} (\sigma_{av}^{\bar{A}_i \pm dA_i})^2 &= \frac{15}{64} \ln^2 \frac{\bar{A}_1 + dA_1}{\bar{A}_1 - dA_1} + \frac{15}{64} \ln^2 \frac{\bar{A}_2 + dA_2}{\bar{A}_2 - dA_2} - \frac{1}{32} \\ &\quad \times \ln \left(\frac{\bar{A}_2 + dA_2}{\bar{A}_2 - dA_2} \frac{\bar{A}_1 + dA_1}{\bar{A}_1 - dA_1} \right). \end{aligned} \quad (2.19)$$

2.6 ‘Man on the Moon’ Experiments of Behavioral Finance

The important insight that we get over and again from behavioral finance is that people are not perfect beings who analyze information and react to it in an optimal manner. Applying this to financial markets, the question is, however, how far do such insights help us understand behavioral impacts on the process of price formation? The reader may have noticed a surprising similarity between the way decision-making is considered in the rational expectations approach discussed in Chap. 1 and the way we have so far discussed behavioral finance in this chapter. In both cases one considers how an individual investor would react to new incoming information, but always in a context where the person is completely isolated, i.e., not interacting with any outsiders, in order to consider his or her decision.

This is particularly clear in the case of behavioral finance, where decision-making is often presented in the form of a questionnaire: given various conditions and choices, how would a given person react? To carry out such questionnaire experiments, the person could just as well be on the moon as sitting at a computer terminal or standing on a trading floor. The title of this section was meant to stress this fact, but also to point out some of the challenges we still need to tackle. Only then will we get a more realistic description of decision-making and a better understanding of financial markets and the way pricing takes place.

It should be noted that prospect theory is no different with respect to such ‘man on the moon’ experiments: the theory deals with human decision-making ‘in isolation’ and has been corroborated by experiments here on the earth, but could just as well have been checked in a capsule on the moon. The same criticism applies to the theories presented in Sects. 2.6 and 2.7, which also only considered

an averaged individual response to information. But financial markets are not isolated objects, which can be understood by assigning typical human decision-making to a representative (average) agent. In short, both the core assumptions of traditional finance (like the dynamic stochastic general equilibrium models of Sect. 1.1) and the core hypothesis in behavioral finance consider the average *individual* response of humans in a *static* or *quasi-static* setting, whereas price formation in financial markets is clearly a *collective* phenomenon with *dynamic* price formation.

Even more intriguing, due to feedback loops, prices may depend on the past price evolution as well as on future expectations. But description of such a scenario, has so far escaped a firm theoretical understanding. Events such as the plunge taken by the British pound back in the 1980s due to short selling by George Soros, or the sudden daily 28 % correction in the 1987 crash, or even the more recent ‘flash crash’ of 2010, all seem like classical examples where a dynamic and collective framework would be needed to provide an adequate explanation.

2.7 Social Processes Underlying Market Dynamics

Social mechanisms underlying non-rational individual choices have received much less attention than the decision mechanisms operating at the level of the individual [51]. As described above, understanding social mechanisms behind the non-rationality of human choices is crucial to understanding the dynamics of financial markets. The dynamics of financial markets is in fact social dynamics. It is driven by social processes and influenced by social factors.

One of the crucial assumptions of social sciences is that individuals react to their own understanding of events and situations rather than to objective reality [94]. Humans actively construct their own representation of reality using their perception and knowledge, but also by consulting with others, especially when they are uncertain. The information and influence received from others shape the understanding of the individual. An important function of human interactions is to construct social representations [99], which are also called shared reality. Social representation for groups is an analog of cognitive representations for individuals. Social representation is inter-subjective: it is a collective view of the world, a mutual understanding and evaluation of situations, objects, and events. Coherent group action is possible if the members of the group share their understanding of the situation, evaluation of objects and events, goals and action plans.

Shared reality thus provides the platform for common action. Humans often treat social representations as objective reality. In their book *Social Construction of Reality*, Berger and Luckmann argue that almost all the social institutions (like, for example, marriage) are created in social interactions and need to be supported by social interactions if they are to exist [17]. In this view, the value of different stocks in financial markets can be understood as a social representation created in interactions among market participants. The value can be negotiated directly in

discussions among investors and their advisors or indirectly by observing others buying and selling financial products, while their actions are reflected in the movements of the stock prices.

Shared reality often takes the form of shared narratives [34]. Narration may be understood as a ‘meta-code’ on the basis of which ‘transcultural messages about the nature of a shared reality can be transmitted’. People tend to store and communicate their knowledge in the form of narratives [15]. In psychology, it has been proposed that scripts, or schematic stories, are the natural representation of knowledge [121]. Narratives are stories that have a beginning, a body, and an end. They have a plot. Narratives describe temporal sequences of actions performed by agents and events and the consequences of actions. Agents have roles, try to achieve goals, and have intentions and plans. Narratives are the instruments by which individuals understand the world in which they live. They link actions to their consequences, they tell agents how to structure their experience and how to evaluate other agents, events, and objects. Narratives are also the basis of prediction. They tell us about the consequences of actions and about typical courses of events. Narratives also tell individuals how others are likely to behave. In terms of narratives, individuals convey economic knowledge, for example, about how economic growth occurs [117].

In a society many narratives usually coexist. Individuals’ interpretation of reality and their expectations depend on the narrative they adopt. Individuals tell stories to others in the social structure they belong to. A narrative is more likely to be adopted if those the individual interacts with also adopt it, or if it reflects a narrative scheme that already exists in society. People attach their personal narratives to the narratives of the groups they belong to. They also change existing narratives and create new ones, usually in a social process, in interaction with others.

Individuals and organizations that are highly placed in power and authority create narratives purposefully in a top–down process. For example, national banks may offer an official interpretation of the financial situation. The intention of such narratives is to impose an interpretation of events in a financial market and to control the market dynamics. Interacting individuals also socially construct narratives in a bottom–up process on the basis of personal experiences [113]. Narratives shared by a social group are created by integration of stories describing the individual experiences of the actors. Shared narratives allow individuals to establish what is common in their experience, find coherent explanations of observed events, and coordinate their decisions and actions. Individuals also construct personal narratives that instantiate common group narratives.

As we shall argue in Chaps. 3 and 8, financial markets are in fact complex systems where the global dynamics emerges from interactions among individuals. Different types of interactions drive the dynamics of the system. Individuals try to understand the situation and the trends in financial markets. Both their attempts to understand and their decisions collectively create an emergent shared reality which is also called a social representation. This shared reality implies the commonality of the individuals’ inner states, views, opinions, and evaluations. Achieving such commonality is motivated, i.e., individuals want to achieve a view of reality that is

congruent with the views of others. It also involves the experience of a successful connection to other people's inner states [42].

The agreed-upon facts, their interpretations, and relations between those facts, as well as their evaluation and emotional connotations, become the basis for investment decisions in financial markets. Because they are shared by a number of individuals, they result in coordinated decisions by a group of people. The communications concern both the global understanding of the market situation and trends and interpretations of recent events, such as rapid changes in indexes, or news published in print or online.

If the decisions of large enough groups of people become synchronized they can influence the dynamics of financial markets. Individual, independent decisions of investors result in a Brownian (i.e., random) motion of stock prices, where small changes are prevalent and large changes are relatively rare. Coordinated decisions stemming from communication among larger groups of people result in price dynamics that resembles Levy flights, a form of dynamics in which the direction of changes is random, but the size of changes varies with each step. This makes very large changes relatively frequent. The dynamics of prices observed in real markets lies somewhere between Brownian motion and Levy steps, which suggests that the decisions of individual investors result partly from communication with others, partly from independent decision-making [31].

Social representation of financial markets is created through different types of communication between heterogeneous agents. The dynamics of financial markets depends to a large extent on the decisions of the largest players, mainly governments, who decide in which currencies to buy and in which to sell. To minimize the impact of these high-volume transactions on financial markets, they are spread over longer time spans. Other high volume players are banks. When the experts working for banks think they know which currency will go up and which will go down, these banks engage in trading. The volume of trading by individual investors is usually much smaller. The decisions of an individual investor usually have a negligible effect on the market. If, however, a larger number of investors synchronize in their decisions, they can influence the market in a significant way. Social representation of the market, also called shared reality, provides one of the main ways in which a large number of investors are likely to synchronize.

Investors communicate with other investors and with financial advisors. Advisors also frequently communicate with other advisors to gain better understanding of what is currently happening in financial markets. They also communicate with other people in their social networks such as their families and friends. Some of the communications are face-to-face discussions in dyads or groups, and such interactions are also often mediated by telephone or Internet. An important medium for the creation of shared reality is Internet discussion groups. Here both advisors and investors can interact with strangers, and this extends the creation of a shared reality. Media, in particular specialized papers and juveniles, also play important roles in the creation of shared reality.

The meso-level also plays an important role in creating the shared reality of financial markets. The social system deciding about the creation of shared reality

has a structure. There are many entities between the level of individuals and the whole system. These entities may have formal or informal structure. Circles of friends, discussion groups, alliances, etc., underlie the formation of clusters in the social network. Individuals who belong to the same cluster will usually make the same decisions, and their actions are likely to be coordinated. Understanding the meso-level of financial markets may represent a significant step in understanding the dynamics of the markets.

Survey research in Poland with 30 investors and 30 advisors gives a more detailed view of the role of communication in market decision-making. In this study, 57 % of respondents answered that they sometimes make decisions on the basis of information from others, and 10 % responded that they do so often or in every case. Communication is especially important when confronted with novelty: 93 % of respondents answered that, at least sometimes, they seek information from others when something new is happening in the market. Social communication appears also to be an important way of dealing with contradictory information: 79 % at least sometimes communicate with others to check contradictory information.

Although the frequency of communication does not significantly differ between investors and advisors the information received from others has more impact on decision-making by investors than advisors. Investors more often than advisors admitted that they take into account the information from others when making decisions and more often change their decisions in consequence of information received from others. Public media press, TV, and radio were indicated as the most significant sources of information. Internet portals and direct communication with others were also indicated as important sources. Internet discussion groups were indicated to be a much less important source of information.

Interacting individuals create shared reality in almost any domain of social life. What is specific to shared reality in the market context is that individuals also interact indirectly through indexes. The activity of an individual is reflected in price movements. Other individuals observe these price movements. By reacting to changes in prices with their decisions to buy or sell, individuals further affect prices. Price dynamics of markets is thus driven by massive feedback loops where the decisions of individuals drive price changes that affect the decisions of other individuals.

Strategies of individuals play a crucial role in determining how current patterns of changes affect further changes, i.e., in establishing the nature of the feedback loops in financial markets. Beliefs about market dynamics underlie individual strategies. The strategies of players may directly influence their decisions, or they may be implemented by computer programs that are increasingly used for trading. The observed dynamics of the market may either reflect a dominant strategy among the players, or emerge from a combination of different strategies. Common features of strategies are likely to cause large swings in the market if the pattern of price changes that is present in a large proportion of strategies occurs in the market. This is due to synchronization of the large number of decisions in the market. Synchronization may be especially pronounced when the decisions are made by computer programs that use the same rules. Technical analysis is used in an attempt

to predict price changes following a specific pattern of changes. Interestingly, since technical analysis is used as the basis for making financial decisions, it may become the basis for strategies of individuals and affect the dynamics of the market.

The most dramatic changes in the market are bubbles and crashes. They occur when individual investment strategies form a positive feedback loop: when the decision to buy by some individuals facilitates buying by other individuals, or the decision to sell facilitates the decision by other individuals to sell. In other words, when rising prices result in a robust decision by a large proportion of individuals to buy, or conversely falling prices result in a decision to sell. Such positive feedback loops may be caused either by cognitive mechanisms based on beliefs and mental models concerning market dynamics, or by emotions, where positive emotions caused by rising value lead to the decision to buy, and negative emotions result in the decision to sell.

The positive feedback loop resulting in bubbles and crashes can also be produced by the fact that, in general, people believe that the price of the asset is what other people believe the price is. As some individuals watch, the value grows. This makes others believe that the value grows, which reinforces the growth belief of the original group. From this perspective, the need for closure [78], i.e., the tendency to hold already formed beliefs, even when faced with contradictory facts, is likely to result in especially pronounced bubbles and crashes, because it makes individuals insensitive to smaller changes in market trends which contradict their beliefs about the direction of market dynamics.

Individual investment strategies are likely to be based on the broader schemas held by individuals. Traditionally, higher stock prices are associated with wealth and profit. Falling stock prices are associated with bankruptcy. There is also a widespread belief that in the long run stock prices rise. This general schema, shared by many market players, is consistent with players taking the long position, which results in market growth. But modern financial instruments allow players to profit equally well from rising and falling stock prices. As an increasing number of players, and especially the big players, become aware of the fact that profits can easily be made on falling stocks, and thus increasingly use the short strategy, the shared reality of market players who associated profit with market growth will begin to change. As discussed above, shared beliefs shape market dynamics. If the shared reality of market players no longer associates profits with rising stocks, the general tendency of markets to grow may no longer be valid.

The general idea expressed in this book is that social dynamics drives financial markets. From this perspective, one can also consider systemic risks, i.e., the risk of collapse of the entire financial system. The resilience of the financial system is for a large part due to the relative independence of the different institutions. If some institutions fall, those remaining can assure functionality of the entire system and create conditions for the fallen institutions to recover, or for new financial institutions to be created in their place. It would require a very large portion of financial institutions to fail for the entire financial system to collapse. The situation changes dramatically when the institutions become strongly interdependent, so that the failure of one institution implies the failure of several others. Financial

mechanisms that may underlie this scenario are relatively clear. Since banks borrow money from other banks, the failure of a bank may have a cascade effect. Social mechanisms may have similar effects.

Information spreads through social networks and results in the failure of financial institutions. As discussed by Christakis and Fowler [31], in September 2009, individuals stood in long queues to withdraw their money from Northern Rock, simply because other people were doing the same. As a consequence, the bank had to close its operations for a day and borrow money from the Bank of England. Many other financial institutions were affected. A short scare experienced by a small number of individuals and spread through social networks can cause a widespread panic that affects several institutions.

Globalization of media, the rapid growth of Internet, and progress in mobile communication result in increasing numbers of people sharing the same information, beliefs, and emotions. People are then also more likely to synchronize their decisions. As a result, financial markets lose degrees of freedom. For example, a cascade of failures of several banks may result from a spreading collapse of trust, rather than from lack of financial resources. If the decision-makers in a bank no longer trust other banks that want to borrow money to be able to pay it back, they may refuse to lend, although there may in fact have been sufficient funds.

Such a collapse of trust can spread through social mechanisms. If the shared reality is such that, in general, no one is sure that the financial institutions will return borrowed money, financial activities may slow down to a level that would not be sufficient for the survival of many institutions. In assessing the risk of systemic failure, it is important to take into account both the financial and the social mechanisms that result in a loss of independence of financial institutions and the globalization of the system [90].

An Introduction to Socio-Finance

Vitting Andersen, J.; Nowak, A.

2013, XIII, 185 p. 44 illus., 9 illus. in color., Hardcover

ISBN: 978-3-642-41943-0