4
Experimental Finance

Debrah Meloso and José Penalva

Introduction

If you have been reading the chapters of this book in order, at this point you are immersed in the workings of modeling and economic theory in general. We ask you to take one step back and picture a financial market. Think of traders sitting in front of computer terminals analyzing data and fancy charts; think of the million things they need to take into account, the complexity of it all; think of the thousands of transactions and the huge amounts of money changing hands every day, every second! And think that at the end of those trades there are people reacting to information and trying to deal with that complexity, with *risks*.

You have seen how people deal with risk before, back in Chapter 2. But now you will see risks as they relate to money, to the flows of capital that originate with firms' investment needs and consumers' borrowing requirements, and that are transformed through financial markets into complex structures that fluctuate and mutate as they change hands across the globe.

Let's give a concrete example: Jim has an excess of cash from an inheritance. He faces a great deal of risk as he evaluates the things he can buy with it, not just today, but over the course of his lifetime (let's not even talk about inflation). Jim considers different options:

- Investing in the high-tech start-up of a friend, Michael.
- Lending money to the local grocer around the corner, who has successfully run her business for over 25 years.
- Stashing the money under his mattress.

P. Branas-Garza et al. (eds.), *Experimental Economics*

The first option links Jim's future consumption with the fortunes of Michael's ideas and management. Jim may also think that having all his cash in Michael's business is not a good idea. So he breaks up the inheritance amongst the different options, building a *portfolio* of investments.

In other words, Jim *diversifies* his portfolio of investments, assuming a little bit of risk (and return) from each of several different investments, rather than a lot of risk from one single investment. Investment portfolios, risk, and diversification are only a few of the basic ideas about finance that you are probably already familiar with. Let us look even further and assume that Jim believes that the three investment options considered so far are too few or unreliable. Jim can then turn to *financial markets*, where he will find lots of people already looking to trade, and where he can transact anonymously knowing that the institution (the market) ensures all deals are honored. The type of markets we usually think about and the type we will study in what follows are *centralized*, *transparent*, and *large*.

Experimental finance and financial markets

As we will see, experimental finance is a huge field of research to which we cannot do full justice in this chapter. We will focus on the study of (competitive) *financial markets*, for they provide the best starting point for delving into experimental finance. The study of financial markets, with its especially solid theoretical constructs, provides a great foundation both for understanding financial issues and the value-added of using experiments.

Markets facilitate capital movements by providing a place to exchange risks. Risks are exchanged in the form of financial *assets*: IBM stock has some risks associated with it, which are different from those of Apple stock. By combining one unit of each, yet another risk profile can be obtained, and so forth. We will study what experiments have to say about how people use markets and the asset prices that come out of them.

Then we will turn to the informational role of financial markets: we will look at what experiments say about the role of markets and prices in transmitting individual information to the whole of the economy, and what this does to people's incentives to acquire that information in the first place.

Finally we conclude with a very quick look at alternative motives for trading assets, the effect of behavioral biases, and a quick overview of other areas of experimental finance.

Risk-sharing and diversification

Returning to competitive financial markets: these are like regular competitive markets (as in Chapter 1) where people act as price takers and where the object of trade is not apples and oranges but promises of future payment (assets). These markets are implemented in laboratory experiments using *double oral auctions* (DOA) or open book markets (see section A in the Appendix), where trading is similar to what you see in regular electronic stock exchanges.

In the lab, these markets are opened for a fixed period of time during which experimental participants buy and sell freely. What participants trade are artificial assets and what we observe is the resulting prices, as in a regular stock market. Additionally, in the lab it is possible to observe whatever assets people hold before, during, and after trade (their initial, intermediate, and *final asset portfolios*). Final portfolios determine the payoff of experimental participants, when they are converted into cash by the experimenter.

It is important to understand why people trade and what the profits from trading are. As we will see, trading is based on differences between agents. We will focus on risk-based differences, but as you saw in Chapter 2, there are other factors that generate differences between agents: the way one weighs probabilities, how one values gains versus losses, how one computes probabilities, even whether in fact one uses probabilities at all. People also trade because they have different information, which raises additional issues which we will address later in this chapter.

There are two primary sources of risk-based differences:

a) Different *initial holdings*. Jim inherits IBM stock while Jenny inherits Apple stock. Jim and Jenny are identical twins, thus their tastes and perception of risk are the same, and their most preferred portfolios should look alike. To get to those portfolios they must trade Apple for IBM stocks.
b) Different *risk preferences*. Jim and Michael hold the same combination of IBM and Apple stock. However, the entrepreneurial Michael

likes to hold more Apple than his more prudent friend Jim. They then benefit from trading and they end up with different investments in stocks (portfolios). Among the factors associated with differences in preferences are: age, health status, family situation, mood, etc.

These differences motivate trade in financial assets. When Jim exchanges exposure to one large risk for lots of small exposures to different risks, we say that he *diversifies*. In order to diversify using financial markets, Jim will have to find another person willing to exchange risky assets with him. In doing this, Jim and other agents involved in financial markets engage in *risk-sharing*. As you have seen in Chapter 2, agents' preference for risks is often to dislike them to a lower or higher degree – they are *risk averse*. Risk-averse agents will like to diversify and thus profit from risk-sharing in financial markets.

Before we discuss our first experiment, you need to meet a few additional important financial concepts. Suppose Jim wants to be able to travel the world with Michael if Michael's high-tech business goes really well. But he is also worried that his grandmother may need home care, as she gets older. To deal with this risk, Jim can invest part of his money in Michael's business, tying his fortune to Michael's, but also put part of his money with the grocer who will be able to pay him back if he needs the money for his grandmother. When Jim does this, he is taking into account that several eventualities or *states of the world* are possible in the future.

In our first example there are four states of the world: Michael's business may thrive, both if Jim's grandmother needs home care and if she doesn't (two states of the world); Michael's business may be a flop, whether Jim's grandmother needs home care or if she doesn't (another two states of the world).

If the assets traded in financial markets suffice for Jim to get any combination of payments across states of the world (in the example we need at least four assets), we say that *markets are complete*. To figure out how the states of the world affect the entire economy, we need to know the *market portfolio*. This is the portfolio of someone who (hypothetically) owns *all* the **risky** assets in the economy. In our example, the market portfolio includes all the shares in Michael's business as well as all shares in other risky stocks and investments.

Instead of investing in risky assets, Jim may want to invest in a *risk-free* asset, which is akin to a really safe bond. Cash acts like a risk-free asset, although usually risk-free assets offer a (very small) return.

Having seen what states of the world and the market portfolio are, you are finally ready to meet *aggregate risk*. This is the risk that cannot be eliminated by spreading one's investments across many assets (*diversification*) or across many people (*risk-sharing*), since it affects the market portfolio itself, which is the most diversified portfolio possible. Formally, there is aggregate risk if the payoff of the market portfolio is different in different states of the world. When we experience a "global" crisis it is because our real world has aggregate risk.

Risk-sharing experiments

We have seen that agents are exposed to great risks and prefer to share them. Markets provide a place where these risks can be traded, where it is possible to share them. There are a number of theoretical models that formalize the risk-sharing function of markets.

Experiments in this section are based on these models.

Consider a simple setting, where all participants have the same information, i.e. no participant has privileged (*private*) information. This allows us to focus on risk-sharing and its implications, and leads us to ask the following questions:

- Will participants trade to change their portfolios (diversify and share risks)?
- Do our models correctly predict the prices of assets and the final portfolios held by participants in the experiments?
- When markets are dynamic, will participants speculate (bet on price movements over time) and will speculation interfere with the profits of risk-sharing (reduce efficiency)?

A static financial market experiment

Bossaerts and Plott (2004) and Bossaerts, Plott and Zame (2007, from now on BPZ) report on a set of experiments aimed at testing the implications of the theoretical equilibrium models of risk-sharing. The main experiment, like all those in this chapter, has the structure presented in Table 4.1.

Table 4.1 Baseline structure of the financial market experiments in this chapter

1. Trading is done via an electronic stock market with an open book market (continuous double oral auction: continuous DOA, see Chapter 1 and Section A in the Appendix).
2. Participants are given cash and initial holdings of a small number of assets that they can trade in this market over a fixed period of time.
3. Participants are not assigned a role as buyer or seller. Instead, they can choose to buy, sell, or hold on to their initial holdings. *Short selling* (selling assets you do not already own) is not permitted.
4. Assets are entitlements to dividends (cash payments). Dividends are paid after trading and their exact value depends on the state of the world.
5. The probabilities associated with states of the world, as well as the relationship between the state of the world and asset dividends, are clearly specified and public information. We call this information the *distribution of dividends*.
6. The experimenter fixes and announces the distribution of dividends prior to the experiment. Participants can make profits from dividends but also from re-trading (buy cheap now and sell at a higher price later). Note that re-trade is a zero-sum game: what one participant wins another loses.
7. There is aggregate risk.

BPZ study one such financial market. In their experiment markets are complete and participants have different (and non-diversified) initial holdings, so we expect them to trade for the purpose of risk-sharing (see Section C in the Appendix for more details). Participants play over several periods in this experiment, and at the beginning of each period participants are "reborn" in an entirely new economy: prices, holdings, and trades that occurred in past periods do not affect initial holdings and the probability distribution of dividends in future periods (see Chapter 1).

The setup is *static*. This does not mean that participants can trade only once. In fact, the continuous DOA allows participants to trade and re-trade at their leisure (and we observe many trades). It is static because:

i. dividends are paid out only once; and,
ii. during a period participants acquire no new information about dividends.

At the end of each trading period, participants receive a payoff determined by their final holdings of assets and the dividends corresponding to the realized state of the world, which is revealed after the trading period ends.

The experiment is designed to capture important elements of the theory of *static symmetric information asset markets*. It is also designed to be able to compute indicators of the presence of risk motives for trade and to discern whether markets are in equilibrium.

We now cover the intuitive components of the theory and equilibrium notions. We consider three closely related models: the Arrow and Debreu model, the Radner model, and the Capital Asset Pricing model. We focus on the indicators and predictions that are tested in the experiments we review in this chapter. The interested reader will find the technical details of equilibrium and related indicators in Section B of the Appendix.

Equilibrium theoretical models considered here capture the following characteristics for markets where agents share risks:

- The world is risky. All risk in the world is contained in the risk of the market portfolio (aggregate risk). The risk of the market portfolio cannot be avoided through diversification.
- All risk-averse agents, trying to reduce their exposure to bad states of the world (states of the world where the market portfolio has low value), will bid up the value of money in such states: in bad states, money is scarce and sources of money in such states are thus expensive.

Although the economy as a whole does poorly in a crisis, some assets may actually do (relatively) well in a crisis (e.g., Wal-Mart or Carriage Services – a funeral services company). Since people are risk averse, if they think a crisis is very likely they will all want to hold assets such as Wal-Mart, thus bidding up their price. These assets are called *countercyclical* because their performance is negatively correlated with that of the market portfolio (when the market portfolio does poorly the countercyclical asset does well, and vice-versa).

To understand the results in BPZ you need a basic understanding of equilibrium concepts and related indicators. The first is the notion of equilibrium used to study the Arrow-Debreu economy (ADE), or the general competitive equilibrium (see Chapter 1).

Arrow-Debreu equilibrium: The ADE treats money in two different states of the world as two different products, like apples and oranges. In the ADE the price of an asset that pays one euro in one state of the world (and nothing otherwise) is equivalent to the price of money in that state of the world. To help visualize the states of the world, consider two states of the world, such as "Rainy" and "Not Rainy."

- If money is very scarce in the state of the world "Rainy" then it will command a high price – just as flawless diamonds are expensive because they are so rare.
- Nevertheless, if the state "Rainy" is very unlikely the price cannot be so high, since even risk-averse agents will pay little attention to a very unlikely (bad) outcome.

Thus, we measure the price of money in a state of the world using the state *price–probability ratio*: that is, the state price divided by the probability of that state. In Arrow-Debreu equilibrium, the price–probability ratio should be highest for states of the world where the market portfolio has the lowest value, and so on.

Radner equilibrium (RadE): RadE is used to study asset prices (instead of the price of money in different states of the world). RadE provides the link between the price of money in each state of the world (ADE) and the prices of assets (e.g., since money is expensive in poor states of the world, countercyclical assets will also be expensive). In this equilibrium, agents trade a given set of available assets, and agents' objectives are to achieve the *best* portfolio they can. The precise meaning of "best" is specific to each agent, but it always relates to the fact that agents are risk averse and need to optimize the trade-off between risk and gain.

A special case of the RadE model – the Capital Asset Pricing Model (CAPM) – assumes that the *best* portfolio is such that the trade-off between the expected value and the variance of the portfolio's payoff is maximized.[1] In equilibrium, the CAPM predicts that all relevant information about an asset's price and, consequently, its *returns* (an asset's *return* in a given state of the world is its dividend in that state divided by its trading price) is captured by the covariance between the asset's returns and the returns of the market portfolio.

More precisely, in the CAPM the equilibrium expected *excess returns* of an asset (the expected difference between the asset's return and that of the risk-free asset) are proportional to this covariance. Also, in the CAPM all agents will end up holding portfolios that look like small replicas of the market portfolio. This will imply that the market portfolio is *mean-variance efficient* in equilibrium, having the highest ratio between the expected excess return of the portfolio and its standard deviation (this ratio is called the *Sharpe ratio*).

In the experiment, BPZ look at several indicators to see if the predictions of the theory hold. In particular they look at:

- state price–probability ratios,
- the Sharpe ratio of the market portfolio, which is compared with the optimal Sharpe ratio (maximal Sharpe ratio that can be obtained given asset prices and payoffs), and
- the relative holdings of risky assets in the final portfolios of the experimental participants, to be compared with the market portfolio.

It is important to note that these indicators are easy to construct in the experiment because the experimenter knows the distribution of dividends and the market portfolio and he/she observes the portfolios of market participants. This is not so in the real world! BPZ's first findings are:

i. Price–probability ratios are ranked as expected: they are highest in the poorest state (X, see the Appendix), lowest in the richest state (Y), and the third state (Z) is between X and Y.
ii. The Sharpe ratio of the market portfolio converges to the optimal Sharpe ratio.

The first finding is consistent with the rank predicted in Arrow-Debreu equilibrium. The second finding (convergence of the Sharpe ratio of the market portfolio to the optimal Sharpe ratio) indicates that prices correspond to those expected in CAPM: such prices suggest that the market portfolio is mean-variance efficient.

Looking back over the questions that were asked about *risk-sharing experiments* at the beginning of the section we find that there is trade (trading volume is high) and prices display some of the properties predicted by the relevant models: state price–probability ratios are ranked as predicted (ADE) and the Sharpe ratio of the market portfolio is close to optimal (CAPM).

However, while prices in BPZ's experiments are consistent with CAPM, participants' final portfolios are not. The CAPM predicts that investors' final portfolios are a combination of the risky part of the market portfolio (scaled down, obviously) with some amount of the risk-free asset.[2] The third result of this experiment is thus:

iii. Participants in the experiment hold risky assets in proportions that are different from the proportions of the market portfolio. However, the mode of participants' holdings follows the market portfolio.

Because prices arise out of the trading which is necessary to attain final portfolios, and prices in the experiment are consistent with the theory, it is particularly surprising that portfolios are not. To explain this, BPZ develop a model where mean-variance preferences are only an approximation of investors' true preferences (see Bossaerts, Plott, and Zame, 2007, for a complete, but very technical description). In this new model, prices and the efficiency of the market portfolio are the same as in the regular CAPM, but equilibrium asset holdings are only on average proportional to the market portfolio.

Using this model, BPZ go back to the data and find that it is fully consistent with the new model.[3] This is a good example of the feedback that can arise between experimental research and the development of theory.

Replication

The BPZ experiment has been replicated with changes in the magnitude of aggregate risk, the type of asset correlations and the composition of the participant pool, and the same results hold. But you run into problems if there are too few participants. Nevertheless, with just three or four traded assets, price convergence is already fast and stable with as few as 20 participants.

Static vs. dynamic financial market experiments

The BPZ experiment is relatively recent and was run after a long series of experiments on financial markets with opposite conclusions, many of which were influenced by the seminal work of Smith, Suchanek and Williams (1988, from now on SSW; in Section D of the Appendix the reader can find a detailed description of the experiment). SSW find that prices diverge from their theoretical ("*fundamental value*") levels – they find *price bubbles*.

The study of financial price bubbles is particularly relevant and interesting, so we will now look at some of the differences between the experiments (SSW vs. BPZ) to understand when bubbles may arise. There are three main differences:

1. The first and most evident difference is that SSW has a *dynamic* setup.
2. In the SSW experiments there is a single long-lived risky asset and cash, while in BPZ there are *several* risky assets.
3. Markets in SSW are not complete.

Consider this step-by-step. The SSW experiment has a *dynamic* setup: that is, participants and their asset holdings "live" for 15 *interconnected* periods. Periods are connected because the final holdings of one period (period *t*) are the initial holdings for the following period (period *t+1*). Also, in SSW an asset is a promise of payment at the end of *every* future period. That is, a participant that starts with and holds an asset during all 15 periods will receive 15 payments. A participant that buys the asset in period 3, buys the right to receive 13 dividend payments (two payments have already been made) plus the right of reselling the asset at any point in the future.

For a similar reason (four possible dividend payments in each out of 15 periods), markets are not complete: there are only two assets that can be traded while there are four *states of the world* in every period.[4]

In the SSW experiments there are two assets: cash and a *single* risky asset that lasts all periods. At the end of every period the risky asset has a 0.25% probability of paying one out of four possible dividends. The stream of dividends is independent and identically distributed (i.i.d.), meaning that the realization of a dividend in one period does not give any new information about the distribution of dividends at the end periods that follow. All this is public information.

In the BPZ experiments, agents hold and trade several *different* risky assets and by re-combining risky assets they can increase the efficiency of their portfolio, increasing the expected gain and reducing its variance. For example, Jim may have an initial portfolio containing only IBM stock while Jenny initially has only Apple stock. Even if Jim and Jenny were equally risk averse, they could both improve the efficiency of their portfolio by acquiring some of the stock held by the other. In that way, their final portfolios would imitate the market portfolio, which – as you have already deduced – is composed of both stock in IBM and Apple. On the contrary, in the SSW setting, with only one risky asset, participants' initial holdings are necessarily a fraction of the market portfolio (why?). The only thing left for agents to do is to decrease or increase the fraction of their total portfolio that is made up of the risky asset vis-à-vis the risk-free asset. If all participants are risk averse, this means that the more risk-averse participants will have to find the less risk-averse ones in order to achieve mutually advantageous trade. We therefore say that SSW provides *weak* reasons for portfolio rebalancing when compared to the BPZ multiple assets setup.

Given the differences we have mentioned, what would you expect to happen in the setting of SSW? Here is what happens.

Consider the *fundamental value* of the risky asset in SSW (see Section D in the Appendix). As time passes there are fewer periods (and payments) until the end of the experiment. In addition, there is no, and there will not be any, new information about dividends. Thus, the intrinsic value of the asset decreases and we expect prices to follow a similar path. Nonetheless, in the experiment we see the very robust appearance of price *increases* that are later drastically reversed in a crash (a bubble that later explodes). The appearance of bubbles is robust to many variations, including the imposition of a price ceiling or the restriction of participation to participants experienced with bubble markets in the past.

At the time SSW ran their experiment their results were taken to indicate that the positive experimental results found for competitive goods markets did not carry over to financial markets. We now know, as exemplified by BPZ, that this is not the main message of SSW. The main message is in fact that, in asset markets, convergence to equilibrium is very sensitive to the different layers of complexity of these markets. We use the three important differences between BPZ and SSW to explore this main message.

SSW propose, and provide evidence for, the hypothesis that speculation in a complex, dynamic environment (difference No. 1) is the driving force for the appearance of bubbles. The main idea is that, in order for prices to follow a path close to the asset's fundamental value, participants need to understand how prices will behave in the future. The experimental data show that, most probably, participants understand how the fundamental value evolves in time (in fact, prices do crash back to the fundamental value in later periods). The problem is that, even if a participant knows the evolution of the fundamental value, she may think that others don't and, hence, that prices will diverge from this value due to the irrationality of other participants. Knowing this, a participant may participate in trade at the "wrong" (irrational) prices, expecting to gain from re-trade to irrational participants at a later time.

Data on participants' price forecasts show that participants do not base these forecasts only on the (well-understood) fundamental value of the asset. Instead, they try to gage market irrationality and *adapt* their forecasts to past forecast errors, supporting the above hypothesis. In general, this hypothesis is called a failure of *common knowledge of rationality*: even if you understand the fundamental value, you believe that others don't and trade at wrong prices to exploit this belief (see chapters 5 and 9 of Vol. 1). In other words, there is a Pygmalion[5] effect in prices that makes them increase only to break down in a big crash towards the end of the experiment, returning to their "rational" level.

Lei, Noussair and Plott (2001, from now on LNP) provide evidence that, even though the above hypothesis may influence the appearance of bubbles, it is not the main driving force. Instead, it is difference No. 3 that captures the essence of pricing bubbles. Let's see how they show this. LNP replicate SSW in an environment where participants have pre-assigned roles as buyers or sellers and, hence, cannot speculate.

Without speculation there is no profiting from others' irrationality over time and, hence, failure of common knowledge of rationality cannot drive a pricing bubble. Nonetheless, LNP observe bubbles in their setting! From this observation LNP elaborate their hypothesis that bubbles are mainly due to *spurious* trade. This is trade that is not motivated by risk sharing but by boredom or a feeling of duty in the experimental setting.[6] Three observations strongly support this hypothesis:

- Bubbles appear in connection with large (unnecessary) trading volume.
- Trading volume falls and bubbles disappear when participants are experienced or when they are given a second task to undertake while trading.
- The LNP (and SSW) setup provides weak reasons for portfolio rebalancing.

In consequence, according to LNP, if the reasons for trade are not those assumed in the underlying theoretical model, there is no reason to expect prices to be as predicted by the model.

What about difference No. 2 (lack of complete markets in SSW)? Is it important? LNP also address this, by running a variation of the SSW setup where the risky asset pays one out of two equally likely dividends in every period.[7] Bubbles do not disappear, thus strengthening the relevance of difference No. 3.

We can now revisit the questions we posed originally in the *risk-sharing* experiments section, armed with the comparison across the three experiments (BPZ, SSW, LNP).

First, there is substantial trading volume, driven by risk-sharing motives (BPZ), by speculation, or – relevant only in the lab – by boredom (SSW and LNP). When driven by risk-sharing motives (as assumed in theory) this trade exploits gains from trade and leads to efficient outcomes. Thus – related to the second question – prices are consistent with theoretical predictions when the motives for trade are as assumed in the theoretical models. Theory's predictions on participants' holdings are supported to a lesser extent by the experimental data: the predicted correlation between participants' final portfolios and the market portfolio is observed for many participants but not all.

Regarding the third and last question, the results of SSW suggest that when the risk-sharing motive for trade is weak and, instead, the speculation motive takes over, speculation does indeed have a detrimental effect on efficiency. However, LNP show that speculation is certainly not the only thing responsible for the inefficiencies seen in a dynamic setting.

We end this section with a remark on experimental methodology. Notice that the theoretical relation between risk aversion and asset prices comes out strongly, provided that participants have a strong desire to hold final portfolios that efficiently trade off risk and expected return. This desire is hampered experimentally if:

- Participants are given efficient and well-diversified initial endowments, or
- Participants expect most of their payoff to come from initial cash holdings and not from their trading choices.

The above points relate to the *salience* property of experiments. An experiment is *salient* if the relation between experimental payoff and making the "right" choices is very strong. The experiments mentioned here show that salience is as relevant in asset-market experiments as in other experiments and that lack of salience hinders results. Whether an experiment has salient incentives or not can only be understood with the guidance of theory.

Informational asymmetries

A fundamental advantage of markets relative to other ways of organizing economic activity – such as central planning in the now extinct U.S.S.R. – is that they bring out information that would otherwise stay hidden.

Returning to Michael's high-tech start-up, Michael may understand very well the workings of his invention, but he may have very little information on relevant factors for its success such as whether there is a demand for his products. Meanwhile there are other agents, for example his potential consumers, who have much better information.

A centralized and transparent market allows everyone to credibly transmit their information[8] through openly observed prices, by "putting their money where their mouth is." We want to see what financial experiments have to say about this information revelation role of markets. We will now consider experiments where:

- The (prior) distribution of dividends is public information (as before), and so are prices.

- But we add the public information that *some investors may have private (privileged) information*.

In particular, some investors called *insiders* have more precise information about the distribution of dividends. One of the main theoretical notions in such economies is that of *Rational Expectations Equilibrium* (*REE* – see Section E in the Appendix). REE implicitly requires agents to do many complex calculations, as it assumes people use all their information in the best way possible, including the information embedded in prices. Also, REE generates very surprising and counterintuitive paradoxes, which we will now consider. Spoiler alert: a way out of the paradoxes is to use a modified version of REE – "noisy" REE.

Where is the paradox in REE and what does it have to do with insider trading? There are a few paradoxes. Let's start with the intuitive claim that an insider will be able to trade and profit from his privileged information, making *informational rents*. Suppose you are trading in an asset market and you are one of those with private information – an insider, like Gordon Gekko in the film *Wall Street*. If you try to buy, everybody else (paranoid that someone out there knows more than they do) will realize that you know the value is going to be high and the price increases. In fact, the moment you start buying, the price will shoot up, become too high for you to profitably trade, and reveal to uninformed traders that dividends are going to be high, that the price adjusts and there is little further reason for trade.[9]

We have just witnessed two paradoxes:

- Because of market paranoia the *insider* cannot profit from his privileged information (no informational rents), and
- prices reveal all information even though there is basically no trade!

These paradoxes are also important to good firm management. In theory, a bad manager will reduce firm value and the firm's share price will fall. A good manager can then buy the firm, fire the bad manager and increase the share price, which would greatly profit the good manager. But if shareholders realize that the firm is going to be taken over by a better manager, they will not want to sell at the low price. They will not sell until the price reflects the value the new manager is going to produce. But, then the price is too high and the new manager will not buy the firm.

Continuing with the paradoxes in REE, consider what would happen if private information in financial markets were not free but instead needed to be acquired with money or effort. In this case the above paradoxes lead to a third one:

- Without informational rents investors have no incentive to acquire information even if it is very cheap.

Thus, the informational efficiency of REE backfires: since nobody will acquire information for the market to aggregate efficiently, prices will reveal no information.

One may argue that none of the above effects will prevail if investors have more reasons to trade than just information. For instance, if you don't like your initial portfolio, you may trade in spite of the paranoia of markets with asymmetric information, because – at reasonable prices – you will still be interested in changing your exposure to risk (diversifying).

The problem is that REE has the power to eliminate even these motives for trade (and again, eliminate the informative function of the market). For example, if the private information of all insiders completely reveals the true state of the world:

- There will be no risk in a REE and, hence, risk-sharing motives for trade will not survive.

This is our fourth paradox, and a real *Catch 22*![10]

Despite these paradoxes, the movie *Wall Street* was a great success (in its time, 1987 – although you may know the 2010 sequel) not just because of the great acting, but because there are Gekkos out there, making money with private information every day, and they are doing just fine.

So where is the problem in the model and how do we change it to be more realistic? The code word used by economists here is "noise." We now turn to experiments where noise helps us find a way out of these paradoxes and improve the theory.

Almost all static experimental markets with private information follow one basic setup introduced by Plott and Sunder in 1982 (from now on PS). It is very special, because assets have *personalized dividends*: different participants receive different dividends in the same state of the world. This "trick" generates reasons to trade that are immune to all REE paradoxes. Even if the REE reveals the true state of the world, investors with a higher dividend in that state will be willing to buy it from those who have a lower dividend. This is not realistic, but it solves a problem that arises when studying "realistic" REE. In a completely informative REE there is no trade and, therefore, prices are not observed. How can we observe prices if nobody trades? Personalized dividends allow us to bypass the lack of reasons to trade, so that the experimenter can observe prices and so that holdings change in markets with asymmetric information. Moreover, this trick generates precise predictions as to who must hold the assets in REE: those investors who obtain the highest dividend in the state of the world that is revealed in REE.

In their experiments, PS assume markets have two or three states of the world whose probabilities are public information. Insiders are either told what the true state of the world is (*concentrated information*) or (in the three-state setup) they are told one of the two states that will *not* be realized (*dispersed information*). Meanwhile, other uninformed investors still believe that all three states of the world are possible.

Numerous variations on the basic PS setup (for instance, Forsythe and Lundholm, 1990) reveal that you can get convergence of market prices to REE levels, but that this convergence depends on several factors:

- First, concentrated information yields faster convergence to REE than dispersed information.
- Second, with dispersed information and (*ex ante*) incomplete markets, convergence occurs, but to an alternative equilibrium notion – one where investors use only their private information and ignore the information contained in prices.[11]
- Third, personalized dividends affect convergence to REE

Concretely, the third point states that convergence to REE is weaker as we go from (i) all investors have the same distribution of dividends, to (ii) investors know that the dividend distribution of others is one out of a few known options, and finally to (iii) investors know nothing about the dividend distribution of others.

Importantly, experiments reveal that, as predicted by theory, whenever there is convergence to REE, insiders do not make extra profits. We are then left with the question of what happens if information is not free – do agents acquire information?

Costly private information

Sunder (1992), and Angerer, Huber and Kirchler (2009), among others, add an information-buying stage before asset markets open in a PS setup. They ask whether information is acquired and, if so, whether REE ensues.

These experiments reveal that the key is whether the experimenter opts for a *fixed slots* setup or a *fixed price* setup. In the fixed slots setup, participants can competitively buy (via auction) a fixed number of "slots" for people to become *insiders*. In the fixed slots setup, perhaps unsurprisingly, as periods progress, investors notice that gains from private information are very low, so that insider slots are auctioned at ever-lower prices. REE emerges because, regardless of how little insiders pay for information, there is always a fixed number of insiders, so that markets remain informative and efficient.

In the fixed price setup, any number of investors can buy information and become an *insider*, as long as they are willing to pay a fixed (low) price. In the fixed price setup, with gains from information close to zero, the number of information buyers, given a fixed price, should go to zero, and, therefore, the market should become uninformative and inefficient. But in the experiment it doesn't because gains from information are always kept sufficiently high to compensate investors for the price paid for information.

How is this possible? Recall our "spoiler alert": noisy behavior (and noisy REE). In the fixed price setup there is unpredictable variation in the number of insiders and in their ability to act on this information. Experimental results portray this variation. This is *noise* that makes other investors uncertain about whether prices are informative or not. This doubt allows insiders to extract some rents (killing one of the paradoxes), which keeps the information flow in markets.

Thus, we observe prices that reflect private information with noise (where the noise is in the minds of non-insiders, who are uncertain as to how much information may be implicit in prices), and this is captured theoretically by the notion of *noisy rational expectations equilibrium* (NREE).

The above analysis of information aggregation in markets was done with experiments in which all private information was concentrated and, therefore, an investor could either be informed (if he bought *all* information) or uninformed. Other experiments have looked at the paradoxes in a context where investors can acquire information of different quality. The first results in this line of research show that only the insiders with the highest quality information can reap informational rents, while other insiders are worse off than investors who choose to buy no private information whatsoever.

Given the importance of noise to the study of information acquisition and trading, some experimenters have introduced it explicitly. Doing this helps them study how noise affects the trading behavior of investors with different market platforms (*protocols*) – other than the standard electronic stock market (DOA). The corresponding notions of equilibrium are very specific and beyond the scope of this chapter.

Other reasons for trade

As we mentioned in the introduction, there are many reasons to trade in financial markets: differences in aversion to losses, ambiguity, differences in mood, and in computational ability. It is *per se* interesting to wonder how these differences can lead to trade that is meaningful (our first question for all experimental financial markets). Even more interesting is to wonder whether, as in REE, the reasons for trade are "eaten" away in equilibrium. Do behavioral biases and computational limitations wash out? Do we observe efficient prices and (portfolio) holdings in equilibrium?

Camerer (1987) addressed this question for biases related to Bayesian updating (see chapter 2 of Vol 1) of public information in dynamic markets. Surprisingly, he found that individual difficulties with Bayesian updating rarely appear in prices (see Section F in the Appendix for the precise updating problem) – just as information differences disappear in REE.

As for ambiguity aversion we find that, on the one hand, individual biases related to risk measurement (e.g., Bayesian updating biases) transform into ambiguity aversion in a market context. On the other hand, ambiguity aversion may not wash away in equilibrium, neither theoretically nor experimentally (Asparouhova et al. (2015); Bossaerts et al. (2010); Section F in the Appendix). A theoretical overview of the issues that may drive the market effect of biases in (financial and non-financial) markets is given in Fehr and Tyran (2005). In particular they find that the key is whether the presence of biases drives non-biased agents to either simulate them (so that biases will be reflected, even exaggerated in prices), or exploit them (eliminating the bias in prices).

Real-world markets are of course complex brews, in which all imaginable motives for trade converge. The power of experiments is not that they realistically replicate these markets, but rather that they isolate the different forces involved in financial markets. In this way, each force's relevance and the quality of models based on them can be understood. What comes out is a better set of lenses with which to look at real-world markets.

Conclusions

We conclude by referring to areas of *Finance* that we have left out of this chapter (see Figure 4.1). We have focused our attention on a couple of aspects related to financial markets and the prices that emerge (*asset pricing*). But finance is more than a bunch of highly paid traders staring at computer screens. Finance encompasses *the study of all existing and potential mechanisms that economic agents use to raise and allocate capital*. Individuals and firms must decide how much they will spend and where they will obtain the funds to finance their expenses and investment projects; similarly, they must decide how to save to "make their money grow" or protect themselves against an uncertain future.

Financial markets describe the interaction between those needing finance and those wanting to invest. But investment and borrowing is not just about share prices and diversification.

Other sub-fields of finance study how financial markets or particular asset structures may emerge, how those markets are organized, and the purpose of the variety of financial actors we observe in financial markets (financial engineering and market microstructure). Most actors are intermediaries of some sort, agents that intermediate between those who can use money for different enterprises and those who have that money (see Section G in the Appendix). Among the many such actors we find banks, insurance companies, investment banks, funds (investment, hedge, and mutual funds), market makers,

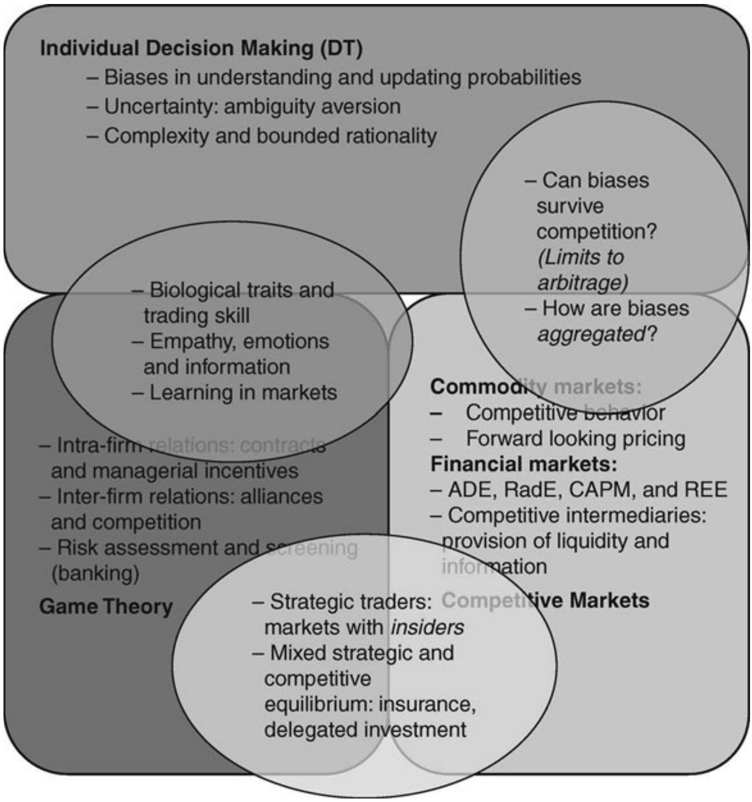

Figure 4.1 summarizes the topics of interest, showing for each case how they overlap with topics covered in other chapters in this book (sometimes they overlap with more than one topic at a time). As you can see, the field of Finance is large! And there exists experimental research relevant to most subfields of finance.

brokers, and (now) trading algorithms. A great deal of effort and research is dedicated to how those markets should function and be regulated (especially in banking and insurance).

Last, but most definitely not least, the decisions firms make about how to finance their expenses and invest their excess cash flows have implications for firm structure and are in turn affected by firm organization. For example, a firm that decides to have its shares publicly traded will often be "owned" by persons who may have very little idea which firm's stock is in their portfolio, let alone how a firm should be managed! Hence, the study of a firm's finances is also concerned with the relationship between stockholders and managers, the internal organization of the firm, and even with the firm's relations with its competitors and allies. These are all participants studied in *Corporate Finance*.

Appendix

A. The open book or double oral auction (DOA)

As discussed in Chapter 1, this is a trading protocol where all participants are allowed to buy or sell assets. Participants trade directly with each other by submitting orders, which are matched to each other through an "open book." An open book is a book of orders listing all the trading orders available for execution/trade, organized by time of arrival and execution price, and visible to all participants. There are two basic types of orders: limit orders and market orders.

Limit order

A limit order to buy specifies an amount and a maximum price a trader is willing to pay to buy that amount of the asset, e.g. "buy 1,000 shares of Google at £625 per share." A limit order to sell specifies an amount and a minimum price a trader is willing to accept to sell that amount of the asset. Limit orders that are not immediately executed build the order book, where they are listed while they wait to be executed.

Market order

The counterpart to a limit order. It asks to buy (or sell) a number of shares at the best possible price, e.g. "sell 1,500 shares of Google." It executes by matching up with existing limit orders. Since this order immediately transforms into a trade, it does not build the order book.

Bid. An order (limit or market) to buy.

Ask. An order to sell.

Trading protocol

A description of the set of rules that regulate trade. The trading protocol specifies the information displayed to each participant (such as whether they can see only the best price to sell and to buy, or the whole order book, or just the best three orders in the book, etc.) the type of offers participants can make, and the way in which participant behavior is allowed to evolve over time. Here we consider only the double oral auction or open book market, but many protocols are used in the real world and in experiments. An important open question is: Why do we observe so many different protocols?

Technology and experimental DOAs

Vernon Smith ran his first experimental DOA (Smith 1962) by building the book on a chalkboard, and keeping track of executed trades in a hand-written notebook. Smith, Suchanek and Williams (1988) used a version of software introduced by Williams (1980) for electronic trading (PLATO). PLATO was well equipped for trading one asset, less so for multiple or correlated assets. Since then, many new electronic DOAs have been developed, including Plott's MUDA, jMarkets (Advani et al. 2003), and the more recent Flex-e-Markets (http://www.flexemarkets.com). DOAs can also be implemented with the flexible experimental programming software z-Tree (Fischbacher 2007).

B. The static setup – theoretical background for experimental data analysis

There are two dates: today and tomorrow. Today investors trade to get from their initial asset holdings to the best possible final asset holdings – they reorganize their asset portfolios. Tomorrow they find out what the state of the world is and consume the money they obtain from the assets holdings in their portfolio.

Arrow-Debreu equilibrium

In an Arrow-Debreu economy, investors can trade Arrow-Debreu (AD) securities. There is a single commodity ("money"), and one AD security for each state of the world – it pays one euro in that state and zero in all other states. So, for example, suppose there are two states of the world: Rainy and Sunny. The AD economy has two securities: holding one unit of the Rainy security provides one euro if it rains and nothing otherwise; holding one unit of security Sunny provides one euro if it is sunny and nothing otherwise.

An Arrow-Debreu equilibrium, ADE, specifies prices of AD securities ($p_1, \ldots, p_s$), and holdings of these securities for each investor, such that: i) given prices, each investor maximizes the expected utility of money in each state of the world, and ii) demand for money equals supply of money in each state of the world.

Interpretation
The study of AD economies isolates the effect that preferences (especially risk aversion) will have on all asset prices. The price of an AD security is the value (today) of guaranteeing one euro in a given (future) state, relative to guaranteeing one euro in another (future) state of the world.

Ranking of state price probability ratios
If agents are all risk averse, it is more valuable to guarantee one euro in a poor state of the world (where the supply of money is small) than in a rich one (where there is lots of money). This means that prices of AD securities divided by the probability of each state of the world, satisfy:

$$\frac{p_s}{\pi_s} > \frac{p_{s'}}{\pi_{s'}} \Leftrightarrow W_s < W_{s'}$$

where s and s' are two states of the world, W is total amount of "money" available, p stands for price, and π is probability.

Radner equilibrium

In a static *Radner* economy investors can trade any number and type of assets. As in the Arrow-Debreu equilibrium, investors also maximize the expected utility of money in future states of the world, but do so indirectly, by constructing portfolios of assets whose payoffs will determine money received in future states. The difference is that assets are more complex and generally they are not AD securities. In this economy, an asset is more valuable if it provides investors with a useful distribution of payments in future states.

A Radner equilibrium (RadE), specifies prices of traded assets ($q_1, \ldots, q_K$), holdings of assets and final consumption (of money) in each state of the world for each investor, such that: i) given prices and asset payoffs, each investor maximizes the expected utility of money in each state of the world, which he obtains from his assets, and ii) demand of each asset equals its supply (sum of initial holdings).

Pricing kernel

If the economy has AD securities, their prices can be used to price all traded assets. That's why these prices are also called the "pricing kernel." In a general RadE the price of an asset is the expected (discounted) value of its payoffs. The expected value is calculated using probabilities that are proportional to the price of (hypothetical) AD securities. Therefore, if we assume a discount factor equal to one, in a RadE with implied AD prices ($p_1, \ldots, p_s$ normalized so they add up to one), an asset with payoff vector $D = (D_1, \ldots, D_s)$ must have the following price q:

$$q = p_1 \times D_1 + p_2 \times D_2 + \cdots + p_s \times D_s$$

In complete markets you can obtain the implied AD prices in this equation from observed asset prices. This implied pricing kernel must satisfy the same ranking property as prices in an ADE (ranking of state price probability ratios). This is a result that is used for data analysis in experimental financial markets.

CAPM

The Capital Asset Pricing Model, CAPM, is a special RadE where investors only care about the *mean* and the *variance* of the distribution of money in future states of the world: that is, investors care about the *mean-variance efficiency* of their portfolios. For each state of the world, define an asset's *return* as its future payoff (dividend) divided by its (current) price. A portfolio is mean-variance efficient if no other portfolio can deliver the same mean future return with a lower variance of returns.

An asset whose return is always the same, regardless of the state of the world, is called a risk-free asset and its return, the risk-free return. The Sharpe ratio of a portfolio is the ratio of the difference of its mean return from the risk-free return, divided by the standard deviation of its returns. It is used to measure mean-variance efficiency. In every CAPM economy, we can compute the maximal achievable Sharpe ratio, which is:

$$\sqrt{(R - R_F)^T \, \Sigma^{-1} \, (R - R_F)}$$

where R is the vector of mean returns of all assets, and Σ is the matrix of asset return covariances. R_F is the risk-free return rate.

In a CAPM equilibrium the market portfolio is mean-variance efficient. Hence, its Sharpe ratio equals the above maximal Sharpe ratio. This is an important measure of convergence to equilibrium in the experiments we study.

C. Experiments

Bossaerts, Plott and Zame, 2007

The experiment encompasses nine sessions. The examples below are computed with parameters of period eight of session 011126: 36 participants, 18 of type I and 18 of type II. There are three states of the world (X, Y, and Z, with probabilities 0.46, 0.27, and 0.27 respectively). Two risky assets, one risk-free asset (*Notes*) and cash.

Table 4A.1 Dividend distribution of assets in francs (F, experimental currency)

	State		
	X	**Y**	**Z**
Asset A	170	370	150
Asset B	160	190	250
Letters	100	100	100

Table 4A.2 The initial asset endowments

	A	**B**	**Letters**	**Cash(F)**
Type I	5	4	–22	400
Type II	2	8	–23.1	400

From Table 4A.2 we learn that the per-capita *market portfolio* is given by 3.5 units of security A, and six units of B. No single participant holds the market portfolio; hence, participants are differently exposed to risk – there is *idiosyncratic risk*.

We combine Tables 4A.1 and 4A.2 to construct the distribution of money across states of the world implied by the market portfolio (Table 4A.3). Clearly there is a lot of aggregate risk. State X, with

Table 4A.3 Earnings of the market portfolio in each state of nature

	The market portfolio
X	170 × 3,5 + 160 × 6 = **1.555**
Y	370 × 3,5 + 190 × 6 = **2.435**
Z	150 × 3,5 + 250 × 6 = **2.025**

the lowest payoff, is a *crisis*, while Y is a state of bonanza and Z is stuck in the middle. (How should state price probability ratios be ranked?)

We give an example of how the state price–probability ratios of Figure 4A.1 are computed using period 8 of session 011126. We use the average trading prices of the last ten trades to compute state

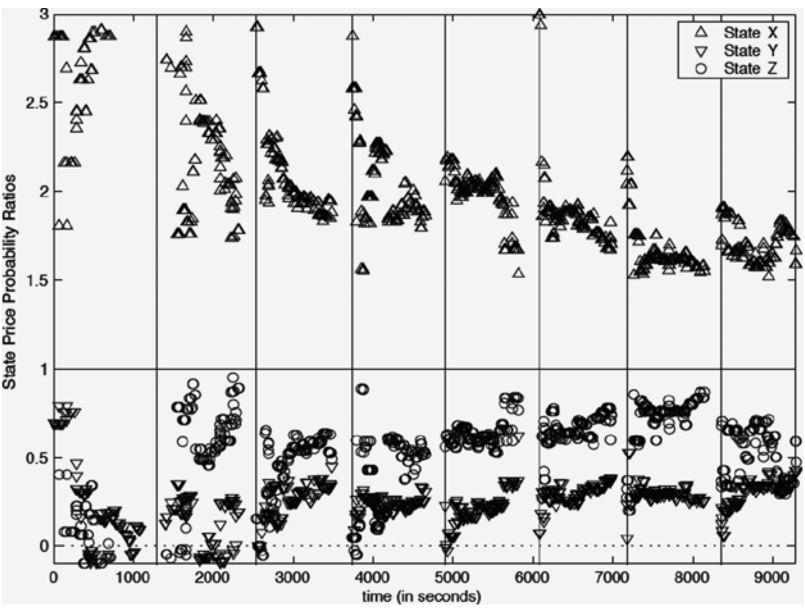

Figure 4A.1 State price–probability ratios

The price–probability ratios are ordered in the direction of the theoretically predicted ranking

Source: Bossaerts and Plott (2004)

price probability ratios: $q_A = 190F$ (price of security A), $q_B = 175F$, and $q_N = 100F$. Remember (Section B) that the price of security A in a RadE satisfies:

$$q_A = p_X \times D_{AX} + p_Y \times D_{AY} + p_Z \times D_{AZ}$$

where p_x is the state price of state X and so on, and D_{AX} is the payoff of security A in state X and so on. This yields a system of equations (in Table 4A.4) that we can solve to find the values of $p_x = 0.76$, $p_y = 0.11$ and $p_Z = 0.13$. We find the state price–probability ratios, p_x, by further dividing state prices by state probabilities, as computed in Table 4A.4. The state price–probability ratios are:

Table 4A.4 State price-probability ratios

• Px = 0,76/0,46 = 1,65,	190 = 170Px+370Py+150Pz
• Py = 0,11/0,27 = 0,41,	175 = 160Px+190Py+250Pz
• Pz = 0,13/0,27 = 0,48.	100 = 100(Px+Py+Pz).

Again, using data from period 8 of session 011126, we give an example of how to compute the Sharpe ratio differences represented in Figure 4A.2. First, compute the securities' returns using final prices: divide the first row of Table 4A.1 by $q_A = 190$ to obtain A's returns in each state of the world. Do the same for the other two securities. Then, use probabilities to obtain mean returns: $R_A = 1.15$, $R_B = 1.1, R_L = 1$, the variances and the covariance: $\sigma_A^2 = 0.24, \sigma_B^2 = 0.045$, and $\sigma_{AB} = -0.013$. Then, $\Sigma = \begin{bmatrix} 0.24 & -0.013 \\ -0.013 & 0.045 \end{bmatrix}$, and the optimal Sharpe ratio is 0.598 (using the formula in Section B). To compute the *Sharpe ratio* of the market portfolio we first compute its price: $p_m = 3.5 \times 190 + 6 \times 175 = 1715$. We use p_m and the dividends in Table 4A.3 to obtain the market portfolio returns: 0.91 (X), 1.42 (Y), and 1.18 (Z). Then compute the mean, the variance, and finally the Sharpe ratio of market returns, which is *0.555*. It is very close to the optimal Sharpe ratio!

D. Smith, Suchanek and Williams 1988

The experiment consisted of 27 sessions, all of them slight variations of each other. We report on session (28x; 9): nine experienced

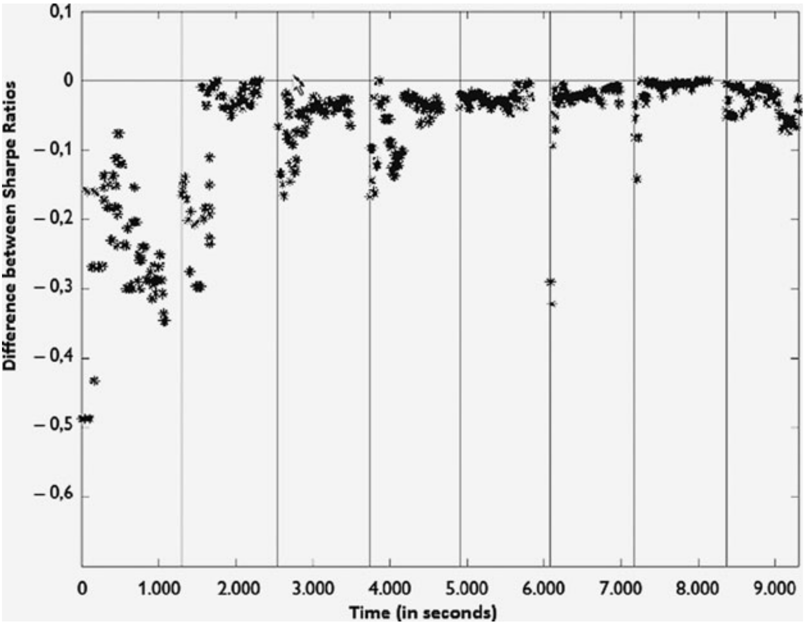

Figure 4A.2 Difference between the *Sharpe ratio* of the market portfolio and the optimal *Sharpe ratio*

Source: Bossaerts and Plott (2004)

participants, three of each type – I, II, or III. Fifteen trading periods with dividend payments at the end of every period, four states of the world per period (states s1, s2, s3, and s4, with probability 0.25 each, independent draws across periods). One risky asset plus cash.

The distribution of dividends given in Table 4A.5 is identical in every period, independent of draws in past periods. Notice that the

Table 4A.5 Dividends of risky assets in each state of nature in a single period (in cents)

	State			
	s1	s2	s3	s4
Dividends	0	8	28	60

Table 4A.6 Initial allocations by type of subject

	Type of subject		
	I	**II**	**III**
Risky	3	2	1
Cash	$2,25	$5,85	$9,45

distribution of *total dividends* changes in time. At time *t=0* there are 4^{15} states of the world, and this number falls at a ratio of four per period as dividend realizations are revealed.

The riskiness of each type's initial endowment is different (Table 4A.6). Notice that – unlike in the multiple-assets case (BPZ), with one single asset – a participant cannot improve his portfolio in both mean and variance. If he increases the mean expected return, the variance will increase, while if he reduces variance he will also reduce the expected return. There is always a trade-off, which is settled depending on the participant's risk aversion.

What about risk aversion? As we saw for the BPZ experiment, an asset's price is not usually equal to its expected payoff. We should therefore question whether the expected sum of dividends is a good reference point to use as the fundamental value of the asset in SSW. The answer is yes! The reason is that since there is only one risky asset, the price of the asset in a market populated with risk-averse participants can never exceed this expected sum of dividend pay-outs. If SSW observed bubble-shaped behavior of trading prices that, nonetheless, remained always below the expected sum of dividend pay-outs, then their results would be questionable. A bubble-shaped evolution of prices where prices move much above the expected sum of dividend pay-outs is incompatible with equilibrium theory and, hence, a real price bubble!

In Figure 4A.3, trading prices are shown as dots connected with solid lines, bids are shown as solid dots and asks are empty dots. The bubble is visible as a large departure upwards from the asset's fundamental value (expected dividend pay-out remaining), shown as straight solid lines. This fundamental value is necessarily decreasing in time, as the asset has ever fewer remaining dividend pay-outs and the expected value of these dividends does not change with time.

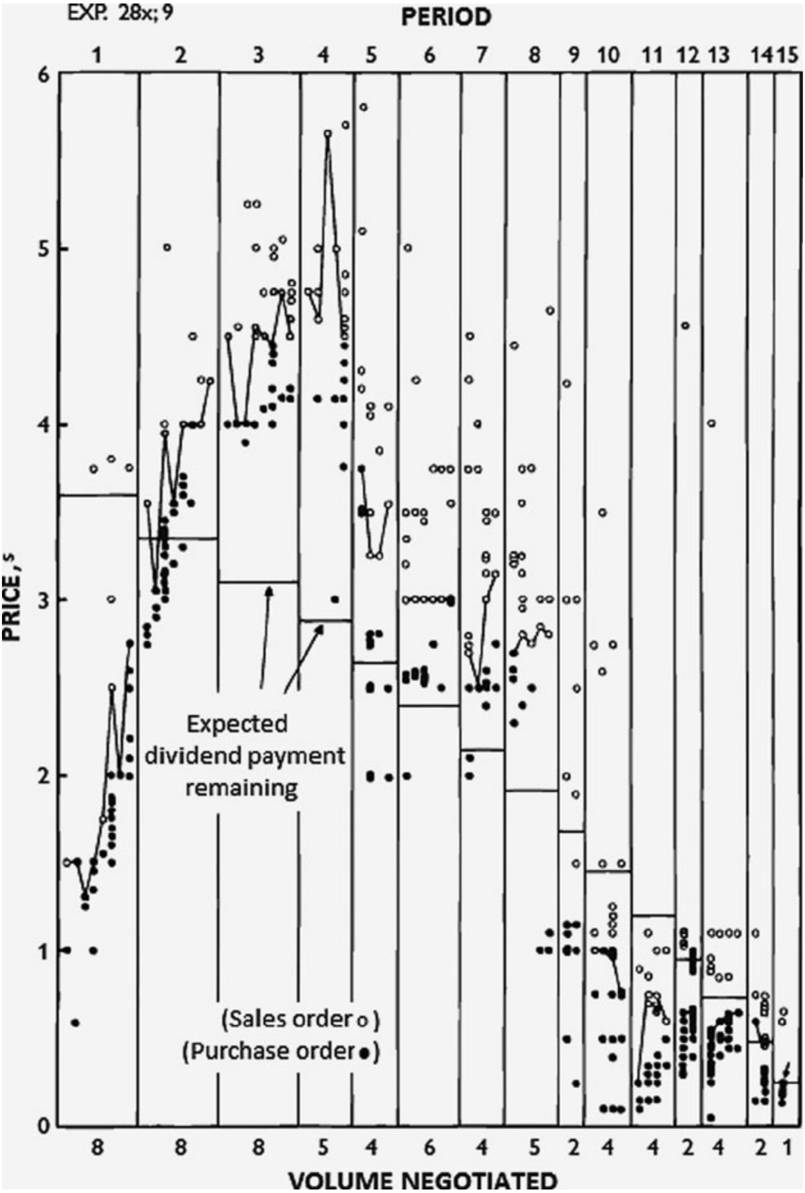

Figure 4A.3 The bubble of session (28x; 9)

Source: Smith, Suchanek, Williams (1988)

We can compute this fundamental value: the expected dividend in a single period is 24 cents. This means that at $t = 0$ the expected dividend pay-out remaining (15 periods: 15 × 24 = 360 cents) is \$3.60. We can compute this fundamental value for every other period by subtracting 24 cents per period. We needn't know the draw of dividend in past periods to compute this! For example, after five periods with their corresponding dividends, there are ten periods remaining and the asset's fundamental value is \$2.40, regardless of what the dividends were in the first five periods.

E. Rational expectation

An investor in a market has *rational expectations* if: when he observes public information he does not regret choices he made in the past – these choices may themselves be part of the public information. In financial markets where there is no private information, *rational expectations* matter in *dynamic markets*. With private information, *rational expectations* matter in both the static and the dynamic case.

In a *rational expectations equilibrium* (REE), we assume all investors can back out private information implicit in prices when they observe those prices. Armed with this power, the REE specifies prices and asset holdings: such that, i) investors maximize expected utility given their beliefs, ii) investor beliefs incorporate all information they can back out from prices, and iii) demand equals supply of assets.

Public information – dynamic

In a setup like that of SSW, demand in one period depends on the distribution of dividends and on the beliefs that investors hold about future prices. If these beliefs are wrong, they will regret the demand they submitted in the past. However, their own demands helped determine the price that ultimately proved them wrong. Only with rational expectations will investors submit demands that generate prices equaling their beliefs. No investor then regrets their choices.

Private information

Example: three states of the world, there is an *insider* who knows only two of these states may occur, the other one will never occur. The outsider submits a demand function based on her belief that all three states are possible. Since she knows there is an insider, once she sees the price, she deduces what information this insider had and regrets her demand.

In a *noisy rational expectations equilibrium* (NREE) the relation between price and private information is probabilistic, so investors cannot perfectly deduce information from prices. Everything we said about REE remains the same here, except that even after seeing the price, investors do not perfectly know what information was in the markets. *Noise* can be motivated in many ways: there are traders that trade "irrationally," blurring prices for everyone else; or, markets are participant to unpredictable supply shocks; or, states of the world do not capture all existing uncertainty.

F. Behavioral bias in markets

Bayesian updating in laboratory financial markets

A large proportion of participants have problems applying the law of probability named Bayesian updating. Consider the following experiment (Camerer 1987): There are two states of the world, X (prob. 0.6) and Y (0.4). There is an urn with three balls in it: In state X this urn contains one red ball and two black balls; in state Y it contains two red balls and one black. At the beginning of the experiment, the experimenter takes out three balls (the *sample*) from an urn (with replacement: takes one out, shows it, and puts it back in the urn) and shows them to everyone. The experiment consists of verifying how people use the information from the balls drawn at the beginning (the sample) when they can trade financial assets. Participants trade a risky asset with dividends of 600*F* in state X and 200*F* in Y.

If the three-ball sample is "two red and one black," what is the probability of state X?

P(X|2 reds) = P(X & 2 reds)/P(2 reds). We use Baye's rule to find that *P(X & 2 reds)*=0.133 and *P(2 reds)*=0.311, implying that *P(X|2 reds) = 0.428*. We can now compute the expected value of the risky asset after seeing a sample with two red balls: *E(asset value|2 reds)* = 0.428x600+0.572x200=328.6.

Theoretical bias – exact representativeness

Notice that a sample of two red balls and one black ball exactly matches the contents of the urn in state Y. A participant displays exact representativeness if she wrongly believes that *P(Y|2 reds)* is close to *one*, when in reality it is 0.572. That is, after observing a sample that matches one of the possible populations exactly, the participant assigns probability close to one to this population. The paper finds support for widespread use of exact representativeness, but it does not show up significantly in asset prices.

Asparouhova et al. (2015) analyze the Monty Hall problem, another bias in Bayesian updating based on a TV show (Google it, it is very cute!). They find that in a market context this bias delivers prices and holdings undistinguishable from ambiguity aversion.

Ambiguity aversion in laboratory financial markets

Ellsberg's one-urn paradox in markets (Ellsberg 1961, Bossaerts et al. 2010): Consider a similar setup as the Bayesian updating ones above. An urn contains red, green, and blue balls. There are nine balls. Of these, three are red and the rest are either green or blue, in unknown proportions. Balls are drawn at the end. Participants can trade three AD securities: security *Red* pays $0.5 if a red ball is drawn and zero otherwise. Similarly for securities *Green* and *Blue*. We call green and blue the *ambiguous* states, since their exact probabilities are unknown.

You are asked whether you prefer security *Red* or *Green*. If you are a median person, you say *Red*. This means you think that *P(state red) = 1/3>P(state green)*. Next, you are asked whether you prefer a portfolio of one *Red* and one *Blue* or a portfolio of one *Green* and one *Blue*. Being a median person, you choose the second portfolio, expressing that you believe *P(state red) + P(state blue) < P(state green) + P(state blue)*, which is inconsistent with your previous choice. This is because the median person is *ambiguity averse*: you dislike betting on objects for which you don't know the exact probabilities. Ambiguity-averse participants, in this experiment, are less reactive than other participants to changes in prices. Their "stubborn" desire to hold a certain type of portfolio, regardless of asset prices, may cause an endogenous change in the market portfolio that is left over for trade among the ambiguity neutral participants. The new (endogenous) market portfolio implies a new ranking of state price–probability ratios. When the authors (Bossaerts et al. 2010) find that state price–probability ratios are ranked according to the endogenous market portfolio, instead of the original one, they find evidence that ambiguity aversion is present and persistent in financial markets.

G. Financial actors

Investors: These financial actors start out holding either cash or a set of assets that they wish to trade into a final portfolio of assets. Based on the distribution of dividends they trade to *control the future payments they will receive*.

Firms: These are complex structures that appear in finance at many levels. They raise capital in financial markets, thus creating *risky assets* for investors to trade; they decide how to finance their operations (debt, capital raising); and they resolve complex contractual problems between owners and employees and across firms.

Managers and stockholders: In the publicly traded firms that are the focus of most of Finance, stockholders own the firm and delegate its operations to a manager. Managers decide what projects to pursue. Their incentives are not always aligned with those of stockholders. The resolution of this conflict is attempted through contracts.

Banks, investment funds, and investment banks: These are all intermediaries that screen and package risky investments. They either bear some of the risk and screen investment opportunities (banks), or they create products that are sold directly (funds) or in competitive markets (investment banks).

Market makers: Intermediaries that trade in financial markets. They accumulate inventories in high supply periods and deplete them in high demand periods. In this way they make markets more liquid. Although they trade in markets they do not have the same (consumption) motivation as investors. They are motivated by the premium that traders are willing to pay in order to trade in a timely manner.

Notes

1. In general, an agent could care about expected value, variance, skewness, kurtosis, and even the shape of the entire distribution of returns. CAPM assumes they only care about expected value (positively) and variance (negatively).
2. An investor with mean-variance preferences will optimally hold a portfolio with maximal Sharpe ratio, and in equilibrium, the market portfolio has maximal Sharpe ratio.
3. Can you think of an explanation as to why experimental asset holdings differ from the theoretical CAPM ones but not prices? To do this exercise you need to understand the contents of Section B in the Appendix very well.

4. We remain superficial about the notion of market completeness since a more detailed discussion would require a long detour from our main topic. It suffices to know that for a market to be complete there must be at least as many traded assets as there are states of the world.
5. The Pygmalion effect comes from a poem by Ovid (*Metamorphosis* X) where Pygmalion the sculptor sculpts a female figure with such beauty (Galatea) that he falls in love with her and treats her like a real woman. Aphrodite has mercy on him and turns the statue into woman.
6. Participants in the laboratory feel "obliged" to do something.
7. We have not talked so far about the meaning of *complete markets* in a dynamic setting. It is more complicated than for static markets. Still, the LNP treatment with only two possible dividends is almost surely a complete markets setup.
8. In Chapter 9, third section, a similar idea is discussed: voting as a mechanism of information aggregation.
9. Symmetrically, if prices don't shoot up in spite of your efforts to buy, you may get suspicious that your private information was inaccurate! This is also an important part of the story that ultimately may paralyze markets with asymmetric information.
10. Catch 22 is a very good anti-war novel by Joseph Heller. Jose and Debrah highly recommend it!
11. *Ex-ante*, that is *before* any private information is transmitted. Private information can reduce the number of relevant states of the world and, thus, transform a previously incomplete market into a complete one.