



UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

The time-variation of volatility and the evolution of expectations

van der Weide, R.

Publication date
2012

[Link to publication](#)

Citation for published version (APA):

van der Weide, R. (2012). *The time-variation of volatility and the evolution of expectations*. [Thesis, fully internal, Universiteit van Amsterdam]. University of Amsterdam.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 1

Introduction and thesis outline

Equity prices, house prices, oil prices, electricity prices, food prices, the price of a currency, they all fluctuate over time. The majority of these different products are traded on active and well-regulated markets where their prices are determined by demand and supply. While demand will largely be driven by consumption, all markets owe part of their liquidity to speculative behavior. The price series are seen to share a number of important empirical features, also known as stylized facts, even though market characteristics may be quite different.

Price bubbles, or boom and bust cycles, denote a prominent example of a stylized fact. During a ‘boom’ prices are seen to grow for an extended period of time, also known as a ‘bull market’. Market participants may believe that the price has risen well above their perception of the fundamental value, but choose not to opt out. Inevitably, prices will come down again, that is, a ‘bear market’ will occur. When they do, prices often come down by as much as they had gained earlier, but considerably faster. This is known as ‘asymmetric volatility’: “Volatility tends to be higher in bear markets” (Engle, 2004).

Another key stylized fact is that the volatility of price returns is universally found to vary over time. Volatility is defined as the standard deviation of the distribution from which the price returns are drawn. While this distribution is not directly observable, such that volatility is strictly speaking invisible, larger returns (in absolute value) are likely to be drawn from distributions with larger standard deviations. Some markets will be more volatile than others, yet most markets see phases of high volatility alternate with phases of low volatility.

This introductory chapter will give a general overview of the themes that are addressed

in this thesis, which can be divided into two parts. Both parts are concerned with capturing the key stylized facts of empirical prices. The first part builds an “econometric model” for the time-variation of multivariate volatility. The objective here is to obtain a good fit of the volatility dynamics so to improve forecasts. In the second part we build “structural economic models” with the objective to gain a better understanding of what economic and behavioral mechanisms are responsible for the observed price and volatility dynamics. To make this opening chapter a coherent story that can be read on its own it embeds the subjects covered by this thesis into a larger review of the literature. The closing section of this introduction provides a focused summary of each of the thesis chapters.

1.1 Price bubbles: Boom and bust cycles

Arguably the first documented example of a price bubble, or the most famous, is the Dutch tulipmania that occurred between 1634 and 1637. In an early example of a futures market, an organized gathering of traders in numerous taverns whose linkage to the physical market is debated, the price of rare tulip bulbs was seen to rise to extraordinarily high levels. At the height of the speculation, early February 1637, prized single bulbs were sold at an equivalent of an average person’s annual income. When prices collapsed, these bulbs reportedly lost around 90 percent of their value in a short period of time. A careful analysis of the tulipmania can be found in Garber (1989, 1990). For a recent discussion offering different historical perspectives on price bubbles, see O’Hara (2008).

The Wall Street crash of 1929 is arguably the most profound stock market crash in the history of the United States. The Dow Jones Industrial Average (DJIA) lost around half of its value in a space of two months. It reached a peak value of 381.17 in September 1929 followed by a low of 198.6 in November 1929. During the boom hundreds of thousands of Americans had entered the stock market, which were in all likelihood drawn by the steady gains the market had accumulated in the years leading up to the crisis. Much of the funds that fueled the speculative bubble was borrowed money. The collapse marked the beginning of the Great Depression, which lasted for more than a decade until the second World War. In 1933 the United States put policies in place that would separate commercial banks from investment banks (the Glass-Steagall Act), designed to prevent history from repeating itself. (Interestingly, governments are considering similar policies today in a response to the global financial crisis of 2008-2011.) See e.g. Shiller (2005) for a discussion of the speculative bubble.

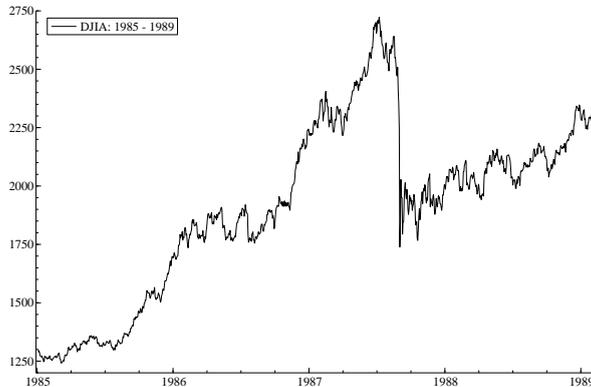


Figure 1.1: *Dow Jones Industrial Average (DJIA): 1985-1989*

Another famous stock market crash occurred about half a century later, the crash of October 1987. Figure 1.1 shows the empirical time-series of the Dow Jones index (a mix of US equity prices) for the period 1985 to 1989. One immediately recognizes the bursting of the bubble on Black Monday, October 19, 1987, after an extended period of steady growth. The index gradually gained about 1000 points over the two year period prior to October 1987, and then lost it all in a matter of days.

A current example, and no less dramatic, is the 2008-2011 global crisis that includes a boom and bust in oil prices, a world food price crisis, and the real estate crisis that fueled a global financial crisis. The latter crises are all inter-connected. Figure 1.2 plots four different empirical price series for the period 2006 to 2009: (1) the Dow Jones Industrial Average, (2) the price of crude oil, (3) the price of Thai export rice, and (4) the Case-Shiller house price index. The time period covers both the boom and the bust. One can clearly see how the global financial crisis inter-links with the crisis in housing, oil and food prices. Note that the decline in house prices preceded, and possibly initiated, the global economic crisis. The price of oil came down after the bubble had burst in financial markets accompanied by economic recession, as this meant a collapse in the demand for oil.

1.2 Price volatility and uncertainty

Figure 1.3 shows the last 30 years of daily (log) price returns for the DJIA (the period 1980 to 2010). It can be seen how volatility varies over time. Periods characterized by large

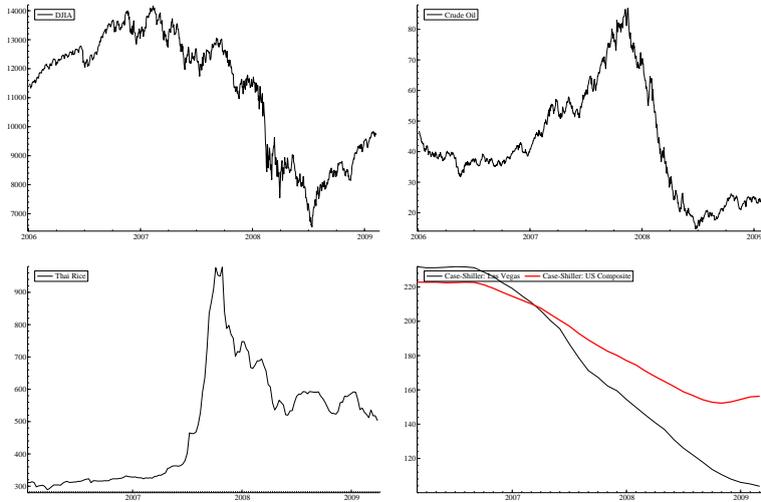


Figure 1.2: Price of DJIA, Crude Oil, Thai Rice, and the Case-Shiller housing Price Index: 2006-2009

(small) price changes may be associated with periods of high (low) volatility. Notice that all return series show a similar pattern: Large (small) price fluctuations today are likely to be followed by large (small) price fluctuations tomorrow, which is an indication that volatility is persistent (today's volatility would be a good predictor for tomorrow's volatility). Phases of high (low) volatility are eventually followed by phases of low (high) volatility, and the cycle repeats itself. This stylized fact is known as 'volatility clustering'. The empirical properties of stock return volatility have been studied by e.g. Anderson (1996), Hamilton and Lin (1996), and Anderson *et al.* (2001). For modeling and forecasting of volatility, see e.g. Anderson *et al.* (2003) and Poon and Granger (2003), and the references therein.

Volatility is associated with uncertainty and risk. An increase in the volatility of an asset's value implies an increase in uncertainty for those who have invested in the asset. A higher volatility means that the investor is at risk of making a higher loss (as well as making a larger profit). The financial sector has made this uncertainty and risk their business by developing increasingly sophisticated ways of measuring and managing risk. It has created a market for derivatives, which denote the primary instruments for mitigating financial risk, also known as 'hedging risk'. Popular examples of derivatives are 'Put' and 'Call' options. The Put option locks-in a minimum value for an asset at which it may be sold in case the actual value falls short of that. The Call option gives an investor looking

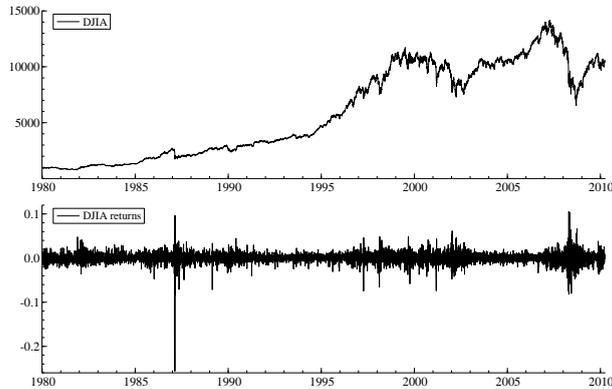


Figure 1.3: *DJIA prices and log returns: 1980-2010*

to buy an asset the guarantee that he/she will not have to pay more than is described in the contract. The prices of these option contracts may be viewed as insurance premiums as it insures the buyer of these contracts against unforeseen price movements.

The uncertainty in future prices and exchange rates brings challenges to many segments of the economy, from large economic agents to small entrepreneurs. Investment banks deal with the uncertain values of financial assets. Real-estate developers look out for changes in land and house prices. A government's monetary policy is in part shaped by expectations about future inflation rates. The profit of import/export companies may be exposed to exchange rate risk.

1.2.1 Macro: Energy price uncertainty

Energy being a primary commodity, large movements in the price of oil and sustained uncertainty in future prices will have significant impacts on the global economy. Bernanke (1983) predicts that optimizing firms will postpone investment decisions when faced with uncertainty about energy prices which will ultimately depress aggregate output. This is confirmed empirically by e.g. Cameron and Schnusenberg (2009), Elder and Serletis (2009, 2010), and Yoon and Ratti (2010). “[V]olatility in oil prices has had a negative and statistically significant effect on several measures of investment, durables consumption and aggregate output” (Elder and Serletis, 2010). Similarly, oil price shocks are found to depress cash flows and lower investments, see e.g. Dotsey and Reid (1992) and Jones and

Kaul (1996). Rising oil prices and heightened uncertainty have also been linked to food price increases (see e.g. Alghalith, 2010). Hamilton (1983) observed that rising oil prices are often followed by major downturns in economic activity.

The large price swings and high volatility in the global market for oil in recent years has renewed the interest into studying the determinants of oil price dynamics. Conjectures concerning the possible determinants include: speculative trading, unexpected contraction in supply (for example when OPEC decides to withhold oil supplies from the market), and global business cycles. The latter would for example attribute the recent surge in prices to the unexpected strong growth in the global economy driven largely by emerging Asia. In a recent empirical study, Kilian and Murphy (2010) find support for the latter, namely that “the sustained run-up in the real price of oil between 2003 and mid-2008 was caused primarily by shifts in the global flow demand for oil”.

Whether speculative behavior also had a part in the rise and fall of oil and food prices of late has recently become a subject of popular debate. The opinions are divided. While it is a tempting conjecture, it is argued that speculators would need to stockpile the commodities for an extended period of time to sustain the rise in prices. The fall in prices would indicate when their stocks are being returned to the market. There is little evidence, however, on the existence of these stockpiles. By the same token, there is tentative evidence that a segment of the oil market in part relies on extrapolating price trends to determine demand, see e.g. ter Ellen and Zwinkels (2010).

It has long been recognized that commodity prices are extremely volatile, see for example Deaton and Laroque (1992): “For some commodities, there have been swings from trough to peak in just a few months. For countries whose export earnings and GNP are dependent on these commodities, such volatility poses major problems both of macroeconomic and microeconomic policy. An understanding of the stochastic processes governing these price movements is essential for macroeconomic management, for national consumption and saving policies, for agricultural pricing policies, and for the design of risk-sharing mechanisms between farmers, resource holders, and government”.

1.2.2 Micro: Firms and farmers

A high price for oil is bad for businesses as it raises costs. It yields high production costs as well as high transport costs. Companies could incur the costs themselves, which would lower their margins and thus curb their profits. Alternatively, companies may opt to recover the extra costs by increasing consumer prices at the risk of lowering the demand

for their produce. For selected products manufacturers will plausibly still see a decline for their produce even if they do not touch the price. Automobiles are a good example, particularly those characterized by excessive fuel consumption. Cameron and Schusenberg (2009) find that when oil prices are high manufacturers of large cars that are not fuel-efficient (e.g. SUVs) are more likely to see their stock prices decline than manufacturers of smaller passenger cars that have a more favourable fuel consumption.

High levels of uncertainty concerning future oil prices arguably also puts pressure on the present value of airline companies. Higher oil prices raises the costs of travel, which airlines cannot always pass on to their customers each time the price of oil goes up, so that it will cut into their profits. For one, they have to compete with other airlines, some of which might be less sensitive to changes in the price of oil, either because they are more fuel efficient or because they were more successful in hedging against the increase in the price of oil. Boswijk and van der Weide (2009) examine the time-varying correlations between changes in the price of oil and changes in the stock prices of selected airlines and transport companies. Upon comparing the airlines, they find that South-West is least sensitive to movements in oil prices, which may be attributed to South-West's active program to hedge fuel prices: "they buy fuel options when the price of oil is believed to be low to hedge against potentially high fuel prices years later. These hedges have for example helped South-West through the Iraq war and hurricane Katrina when oil prices indeed increased significantly. In the third quarter of 2008 ... South-West recorded its first loss in many years, partly because the then noticeable drop in oil prices had rendered the fuel hedging strategy of lesser value" (Boswijk and van der Weide, 2009).

Correlations are found to peak in periods with increased uncertainty in future oil prices; in 2002 with the build-up to the Iraq war, and in 2008 when oil prices reached their peak against the backdrop of a looming financial crisis. "At the outset of the financial crisis, oil prices were still climbing ... Thus when the recession in the United States was already looming, the price of oil was still high ... there was considerable uncertainty at the time whether the recession would bring down oil prices. Much of the (new) demand for oil has come from emerging economies such as China, as well as from Europe. It was uncertain if and how the recession in the United States would carry over to the emerging markets and Europe. In case of a decoupling the demand for oil would stay strong despite the recession originating in the United States. In this time of uncertainty, correlations increased in absolute value. When it then became clear that the recession would be global also oil prices came down, together with stock prices" (Boswijk and van der Weide, 2009).

A farmer would be an example of a smaller entrepreneur for which uncertainty in future prices brings a challenge. “Farmers the world over, in dealing with costs, returns, and risks, are calculating economic agents. Within their small individual, allocative domain they are fine-tuning entrepreneurs” (Nerlove, 1979). For the planning of the crop production, i.e. when deciding on how much to cultivate of each crop, the farmer will be relying on his/her estimates of post-harvest prices. Where farmers will inevitably make errors in the forecasting of agricultural prices, it may be expected that rising uncertainty concerning future prices will put pressure on their profits. “In so far as farmers formulate production plans on the basis of their price expectations, the failure of expected prices to correspond to realized prices will prevent ex-post profits from being maximized. Equally important, errors in price expectations may impair the efficiency of agriculture” (Heady and Kaldor, 1954).

In developing countries, where both infrastructure and markets are not yet fully developed, building expectations of future prices is often hampered by poor access to information. When a market for derivatives and insurance is also missing, farmers will find it difficult to hedge their risks against unforeseen price movements, whether that is due to demand shocks or supply shocks. Yet, also farmers in the developed world, where access to information is more efficient and where derivative markets are available to hedge against price risks, the price fluctuations in the commodity markets still represent a significant difficulty.

In April 2008, the New York times published an article entitled ‘Price Volatility Adds to Worry on U.S. Farms’, for which they interviewed a farmer from Illinois, Fred Grieder, who has been farming on 1500 acres for 30 years. A typical days work may range between 12 and 20 hours, which means long days of plowing, planting, fertilizing, and looking after the crops. “But Mr. Grieders days on the farm ... are getting even longer. He now has to keep a closer eye on the derivatives markets in Chicago, trying to hedge his risks so that he knows how much he will be paid in the future for crops he is planting now. And the financial tools he uses to make such bets are getting more expensive and less reliable ... todays crop prices are not just much higher, they also are much more volatile ... The price swing expected in March for soy beans was three times its monthly average, and the expected volatility in corn prices was twice its monthly average ... there is no question that the grain markets are now experiencing levels of volatility that are running well above the average levels over the last quarter-century ... Mr. Grieders crop insurance premiums rise with the volatility. So does the cost of trading in options, which is the financial tool he has

used to hedge against falling prices ... In any case, at current levels of volatility, options trading becomes riskier, and therefore more expensive, too expensive for many farmers like Mr. Grieder, who now has to hedge with the recently less reliable futures contracts. That exposes him to the risk of having to put up more cash, to maintain his price protection, whenever a weather threat, shipping disruption or a fresh surge of money from Wall Street suddenly pushes up grain prices ... 'If you have got 50,000 bushels hedged and the market moves up 20 cents, that would be a USD10,000 day,' he said. 'If you only had USD10,000 in your margin account, you would have to sit down and write a check ... On an unusual day, he said, he might get four phone calls a day from his broker seeking additional margin ... he sometimes has to rely on his bank to advance him the margin he needs to keep those hedges in place, a worrisome requirement even for a successful farmer in an economy already struggling with a credit squeeze ... Farmers used to leave the market-watching to traders who work for big grain elevator companies. But with some of those companies now refusing to buy crops in advance because hedging has become so expensive and uncertain, farmers have to follow and trade in those markets themselves" (The New York Times, 2008).

1.3 The determinants of price- and volatility dynamics

What makes prices fluctuate to the extent that they do? Can the enormous swings in prices that define a boom and bust cycle be justified by changes in fundamentals, or are there more forces at work? Naturally, markets will re-evaluate the price of an asset as news comes in that concerns the intrinsic value, think of earning announcements, newly found demand, technological discoveries, changes in the price of energy, changes in the interest rate, plans for a merger or acquisition etc. For example, Garber (1989, 1990) hypothesizes that during the boom and bust cycle of the tulipmania (1634-1637) the price of rare bulbs were still in touch with fundamentals. The nature of the process that ultimately transforms rare bulbs into commercially viable tulips, which essentially gives the owner the rights to the production of the new species of tulips, offers a rationale for the dramatic price movements. In fact, similar price swings can still be observed today albeit to a lesser extent due to technological advancements over the years.

There is evidence however that suggests that these news events are not fully able to explain the magnitude of observed price changes in general, i.e. price changes are believed

to be more volatile than changes in the fundamental value. Shiller (1990) refers to this discrepancy as ‘excess volatility’. Also volatility is found to fluctuate, see e.g. Schwert (1989) who addresses the question as to “Why does stock market volatility change over time?”. What makes prices and volatility fluctuate can arguably be traced back to both economic and behavioral factors.

Among the prominent economic determinants of market volatility are the volatility of macroeconomic factors such as GDP growth, industrial production, inflation, short term interest rate, and the exchange rate (see e.g. Engle and Rangel, 2008; and Engle *et al.*, 2009). For some macroeconomic factors levels too are found to play a role, “there is evidence that high inflation and slow growth of output are also positive determinants” (Engle *et al.*, 2009). For agricultural commodities, other economic factors such as the cost of storage may have an impact on price volatility.

Behavioral factors too may shape the dynamics of prices, think of economic agents updating their expectations or adjusting their degree of risk aversion (likely triggered by news events and/or observed price fluctuations). Economic agents determine how much they buy or sell of an asset based on their expectations concerning its future value in an effort to seek a profit or avoid a loss. In a world where prices fluctuate, at a minimum driven by changes in the fundamental value, agents will want to anticipate price movements. This leads them to form expectations about future prices, which they will update as new information (that includes newly observed prices) becomes available. A ‘feedback loop’ emerges where today’s prices shape the expectations that will determine tomorrow’s prices. It can be shown that this expectation formation process will generally create fluctuations of its own. Moreover, it is conceivable that the ‘endogenous dynamics’ will also act to amplify the significance of exogenous shocks to prices. The ways in which expectations shape the dynamics of prices can take different forms. See e.g. Brock and Hommes (1997, 1998), Lux and Marchesi (1999), Evans and Honkapohja (2003), Boswijk *et al.* (2007), Gaunersdorfer *et al.* (2008), Heemeijer *et al.* (2009), Burnside *et al.* (2011), and the references therein.

1.3.1 Rational expectations

The economist will need to make assumptions about the economy, think of market structure, production, depreciations and costs, as well as about the behavioral underpinnings, think of preferences and expectations. Here we focus on the latter, in particular on the expectation formation process. The standard assumption is that agents have ‘rational expectations’. This is often assumed to imply that agents either have perfect foresight (in

deterministic models) or make forecasting errors that are zero on average (in stochastic models). In the original theory of rational expectations put forward by Muth (1961), it is not so much assumed that individuals (firms in Muth's case) were rational, but that the market as a whole, the aggregate, is rational. Individual forecasting errors are assumed to cancel out on average, so that the average individual has perfect foresight. This is weaker than assuming rationality at the individual level. Lucas (1972) denotes an early and well-known example that adopts the rational expectations framework.

In general, however, we have that a collective of agents with heterogeneous expectations behaves differently than a single 'representative agent' holding their average expectations. In other words, Muth (1961) relies heavily on assumptions of linearity. That does not mean that the hypothesis of Muth (1961), that aggregate expectations are rational, is unreasonable. It finds supports in early survey data on expectations, see e.g. Heady and Kaldor (1954): "For example, the average forecast for corn in 1949 differed from the realized price by about 1 per cent. Yet, only 52 per cent of the individual forecasts fell within plus or minus 10 per cent of the realized price".

Over time, rational expectations prevailed as a popular assumption, also when the linearity assumption is not satisfied. This means that in general all agents are assumed to be rational. It largely owes its success to its analytical tractability, and that the alternative is less well-defined, i.e. if expectations are not rational they could be almost anything.

Note that rational expectations are not necessarily inconsistent with price bubbles. It is harder however to rationalize boom and bust cycles, which is what we observe in practice. In a standard rational expectations asset pricing model, prices can either follow the fundamental or grow indefinitely, away from the fundamental (see e.g. Flood and Garber, 1980). For boom and bust cycles to emerge we need prices to switch between the fundamental and the bubble solution. This would arguably require a non-rational intervention; think of a fraction of the market questioning the bubble, fearing it has to stop some day.

In the end rational expectations is a rather strong and thereby an unrealistic assumption. Perfect foresight implies perfect knowledge of the underlying model, which often requires knowledge of the expectations of other market participants. Also, if all agents were rational, we would arguably not see as much trade as we do, see the no trade theorems of e.g. Milgrom and Stokey (1982) and Tirole (1982). Rational expectation models are arguably not well equipped to rationalize the observed data.

1.3.2 Boundedly rational expectations

As the evidence mounted against rational expectations the concept of bounded rationality emerged as a more realistic alternative assumption, see for example the survey by Sargent (1993).¹ “Economics and finance are witnessing an important paradigm shift, from a representative, rational agent approach towards a behavioral, agent-based approach in which markets are populated with boundedly rational, heterogenous agents using rule of thumb strategies [...] A boundedly rational agent forms expectations based upon observable quantities and adapts his forecasting rule as additional observations become available” (Hommes, 2006). The forecasting rules can be as basic as extrapolating trends or as sophisticated as econometric modeling where unknown parameters are identified from the data. Learning can also be accommodated, which may entail that the econometrician updates his/her model parameters as new data becomes available, or he/she might switch to a new model all together if that fits the data better.

Different agents may work with different models, update their models using different criteria and operate at different time horizons (think of long-term versus short-term investors), which introduces natural levels of heterogeneity. Models of expectation formation that go by these rules are deemed an accurate and realistic account of human behavior.

In theory, boundedly rational expectations may in time converge to rational expectations, if the underlying model does not vary over time (or very gradually) and if the model is not too complex, so that a Bayesian econometrician has a fair chance of learning the underlying mechanisms (see e.g. Cyert and DeGroot, 1974). By the same token, this may realistically never happen. Also heterogeneity might steadily decline over time until all agents agree with each other on future prices (in which case trade may cease to exist), but this probably too will not happen.² Realistically, any given market will have heterogeneity in abundance.

Consider the following three different mechanisms by which the expectation formation process may introduce endogenous dynamics in prices: (i) rational herding where agents will follow recent trends in prices even if that means a deviation from the perceived fundamental value; (ii) performance based switching between cost-efficient naive estimates and costly sophisticated estimates; and (iii) changing weights given to prior beliefs relative to observed data. Let us briefly elaborate on each of these mechanisms.

¹Early ideas of bounded rationality can be traced back to Simon (1957), before Muth (1961) launched his concept of rational expectations.

²For evidence that heterogeneity is likely to prevail, see e.g. Lahiri and Sheng (2008).

Rational herding (see e.g. Scharfstein and Stein, 1990; Froot *et al.*, 1992): With this we have in mind a world where agents will follow each other, which may temporarily drive prices away from the fundamental value. One could think of different reasons why such herding behavior may be optimal. It is conceivable that following the crowd may at times lead to more accurate price forecasts. In periods where speculative forces are stronger than those of the ‘underlying economic system’, prices may be expected to diverge from the fundamental. It would then be rational to follow the herd rather than coordinate on fundamentals (see e.g. Froot *et al.*, 1992). It can be shown that a world where agents extrapolate price trends to predict future prices is prone to endogenous price fluctuations (see e.g. Hommes (2006) and the references therein).

Switching between naive and rational expectations (Brock and Hommes, 1997): It does not always pay to invest in costly estimates of future price movements. It only pays when the benefits exceed the costs. With a stylized model Brock and Hommes (1997) illustrate how a calculated switching between cheaply available ‘naive expectations’ and costly ‘rational expectations’ will introduce price dynamics where prices will fluctuate close to equilibrium for an extended period of time, but eventually move away from equilibrium. It is assumed that rational expectations have a stabilizing effect, i.e. when the entire market acts rationally, prices will tend to the equilibrium value. Similarly, naive expectations are assumed to destabilize the market. Consider an initial state where prices fluctuate close to equilibrium. In this case it is not hard to predict tomorrow’s price. Naive expectations will then be just as good as the more sophisticated expectations, but cheaper. As naive expectations are now the cost-efficient option, more and more agents will shy away from investing in expensive estimates. Until almost the entire market is adopting naive expectations, which is when the price dynamics becomes unstable. Prices will start to diverge from equilibrium. As price fluctuations become larger, they also become more unpredictable. At some point the benefits no longer outweigh the costs of acquiring rational expectations, and hence rational expectations will start to attract more agents. Until the market as a whole becomes rational, which is when prices will return to equilibrium, and we are back to where we started.

Choosing between prior beliefs and observed data (Diks and van der Weide, 2011): In this world, agents act as Bayesian statisticians, they learn about the underlying model from observed price data, and combine this with their prior beliefs to obtain estimates of future prices. Analogous to Brock and Hommes (1997) it is assumed that the market is unstable when all agents entirely discard observed prices when building price expectations.

It is also assumed that agents believe in a changing world, i.e. that the model underlying the price dynamics from before the 1987 stock market crash (in a world without internet) is believed to be different from the model underlying today's price dynamics. This means that agents will give more weight to recent price data relative to older price data, and will eventually forget about distant observations. The mechanism works as follows. Consider again the initial state where prices are close to equilibrium. When prices show very little fluctuations, they offer little information. In time, the last period of substantial price variation will gradually be considered less relevant as the underlying model may have changed over time. If observed prices remain uninformative, in the form of little price variation, agents will start to attach more weight to their prior beliefs, and in effect less and less to observed data. Until prevailing market expectations are no longer shaped by observed data, which is when the market becomes unstable, and prices will start to deviate from equilibrium. Prices now fluctuate, making observed price data informative. Soon agents will be trading their priors for data. Then, as the underlying model is becoming known to the learning agents, the market stabilizes, and prices tend to equilibrium.

1.3.3 What do real-life expectations look like?

Models that accommodate a more realistic expectation formation process are more likely to be consistent with empirical data. What makes expectations more realistic? To answer that question, different ways of collecting data on expectations have been explored, think of: (i) surveys that directly ask for expectations held by economic agents such as farmers or financial analysts, (ii) laboratory experiments that allow the researcher to monitor the expectations and actions of human subjects as they interact in an artificial market, and (iii) inferring market expectations from observed price data. (The latter is only a realistic option for selective markets, such as the markets for futures, forward exchange rates, and options.)

In the mid-twentieth century, Heady and Kaldor (1954) conducted a three-year study of expectations held by farmers from the ten southern counties of Iowa in the United States. While farmers and producers in the United States have obviously developed over time, some of the distinctive features of expectation formation may be expected to apply today still. Moreover, some of the insights obtained by studying farmers from the developed world in the 1940s (shortly after the second World War) may well carry over to today's farmers in the developing world. From their interviews they learned that the farmers commonly "start the process of devising expected prices from current prices. The current price was then

adjusted for the expected effects of important supply-and-demand forces. Where farmers possessed little information about these forces, there was a tendency to project either the current or the recent price trend". When it comes to forming expectations, financial specialists are not that different from farmers. Evidence on exchange rate expectations held by financial specialists, collected by different surveys over the years, is summarized in Hommes (2006): "A consistent finding from survey data is that at short horizons investors tend to use extrapolative chartists' trading rules, whereas at longer horizons investors tend to use mean reverting fundamentalists' trading rules".³

Laboratory experiments confirm that humans, when faced with simple decision problems under uncertainty, such as planning production or trading on a market, are more likely to employ simple heuristics than to act as rational agents. For example, the asset pricing experiments by Smith *et al.* (1988) show how price bubbles can be seen to emerge in artificial markets where all market participants have sufficient information to compute the fundamental value of the asset that is traded. This is confirmed by Hommes *et al.* (2005) whose experiments show that persistent price deviations from the fundamental are reinforced by trend following strategies. Subsequent experiments by Dufwenberg *et al.* (2005) and Haruvy *et al.* (2007) however find that these price bubbles tend to shrink or are eliminated all together when experienced traders are included (even if they are outnumbered by inexperienced traders). This tells us that laboratory experiments can be learned which then stabilizes the experimental market. Real markets however are considerably more complex. Waves of new technologies and exogenous shocks mean that real-life conditions are ever changing making them considerably harder to learn. In a recent experiment, Hussam *et al.* (2008) try to mimic these conditions by introducing increased dividend and liquidity uncertainty in which case bubbles are seen to re-emerge even with experienced traders. "[I]ncreased experience in the same environment is the only condition that has reliably eliminated price bubbles" (Hussam *et al.*, 2008). Inexperienced traders plausibly still have a part in nurturing price bubbles (with experienced traders unable to eliminate them). This is confirmed in a recent empirical study by Greenwood and Nagel (2009). Using data on actual mutual fund managers they find that inexperienced managers are more likely to adopt trend-chasing behavior, tend to extrapolate from the limited data they have observed in their career, and are more prone to the type of optimism that fuels bubbles. Consistent with these observations is that at the peak of the technology price bubble, it was the less experienced fund managers that were more heavily invested in the

³See e.g. Frankel and Froot (1987, 1990), Allen and Taylor (1990), Ito (1990), and Taylor and Allen (1992).

riskier technology stocks.

Observed price data too may hold information about expectations. Market prices are made from different ingredients, of which market expectations is one. While expectations are seen to go into market prices, it is not given that one can also get them out. In some cases you can. From option price data for example one can infer market expectations about future financial volatility: “Because option value depends critically on expected future volatility, the volatility expectation of market participants can be recovered by inverting the option pricing formula” (Dumas *et al.*, 1998).⁴

Jackwerth and Rubinstein (1996) and Bates (2000) are among the studies that exploited option price data to learn about the subjective distribution of volatility. They find that the subjective distributions of volatility from before and after the stock market crash of 1987 look very different. There is no evidence, however, that the objective distribution of volatility experienced a similar change. That suggests that market expectations have changed. The way the market now models the volatility process is consistent with the objective distribution of volatility, while volatility expectations from before the crash were clearly misspecified. This shows that market participants are not rational agents, but that they are learning. At times, learning may be triggered by dramatic events.

There is also evidence that economic agents are different, perceive information differently, and form different expectations, as might be expected. In a recent paper, Lahiri and Sheng (2008) note in their opening paragraph that “various survey data on expectations from many countries over [the] last fifty years have produced mounting evidence on substantial interpersonal heterogeneity in how people perceive the current and form inference about the future economic conditions”. Heady and Kaldor (1954), in their study of U.S. farmers in the late 1940s, find that farmers use a variety of models involving various degrees of complexity: “No single procedure was employed by all farmers. Moreover, the same farmer often used more than one procedure, depending upon the amount of information possessed and upon the degree of confidence attached to it”. This shows that economic agents do not only adopt different models, over time they also make calculated switches between models.

⁴Their statement holds true because the option value as a function of volatility is indeed invertible.

1.4 The benefits of having a good model

What are the benefits of modelling price dynamics? That largely depends on whether one is building an economic or an econometric model. Risk measurement, portfolio allocation, derivative pricing all benefit from having a statistical model that accurately describes the time-variation of asset prices including its volatility process. A theoretical model of the economic and behavioral mechanisms underlying the observed price dynamics offers different benefits. It may offer explanations for observed regularities or anomalies, such as the recent price swings in the oil and food markets. By tweaking the models they may be used to study the effects of changes to the economy, think of innovations that alter the cost functions of producers. Similarly, economic models may be adopted to study the implications of policy interventions, think of price subsidies, regulations that alter market trading rules etc.

For example, the hope is that a better understanding of what is underlying the dramatic fluctuations in oil prices will help policy makers to decide between policies designed to curb the fluctuations in the oil market. Under the assumption that speculative trading plays an important role, “there has been considerable political pressure recently to impose regulatory limits on trading in oil futures markets” (Kilian and Murphy, 2010). Alternatively, if “surges in the global business cycle are the chief cause of high oil prices, then efforts aimed at reviving the global economy after the financial crisis are likely to cause the real price of oil to recover as well, creating a policy dilemma” (Kilian and Murphy, 2010).

What does a good model look like? Also this depends on whether one is looking at an economic or an econometric model. For an econometric model all that matters is that it fits the data well and provides good out-of-sample forecasts, so that risks will be accurately measured and derivatives will be accurately priced. An economic model too is judged on how well it is able to describe the observed data, but importantly, the model is also judged on whether it manages to do so by relying on sensible economic and behavioral assumptions.

One could conceivably construct different economic models that are capable of replicating the same price dynamics. When tweaked, however, they may lead to different results. Thus, the relevance of the model as a tool for studying policy interventions relies crucially on whether the model has captured the appropriate underlying mechanisms. The assumptions the model is built on play an important role here. “Evidence from behavioral finance helps us to understand for example that the ... worldwide stock market boom, and then crash after 2000, had its origins in human foibles and arbitrary feedback relations ... The challenge for economists is to make this reality a better part of their models” (Shiller,

2003).

1.4.1 Building economic models

Having many agents with different expectations interacting with each other, however uncomplicated each individual rule, will likely lead to complicated dynamics. There is obviously a trade-off between realism and complexity. A strand of literature that wishes not to compromise on the complexity inherent to more realistic models of expectation formation relies on computers to trace the large numbers of interacting agents. Using computer simulations one can study aggregate outcomes, such as market prices, profits, trade volumes, and levels of heterogeneity without having to impose constraints on the number of trading rules and learning mechanisms (for an overview, see LeBaron, 2000). This is how the Santa Fe artificial stock market came into existence (see e.g. LeBaron *et al.*, 1999). “In the Santa Fe market, the actions of agents, which are based upon their expectations, are explicitly modelled and traced for each individual agent. With computers becoming cheaply available and faster, these artificial markets allow for more and more detailed modelling” (Diks and van der Weide, 2005). The approach also has disadvantages, one is that the mechanisms underlying the aggregate outcomes essentially become a black box (just as with observed data), i.e. “it is not always clear what exactly causes an observed simulation outcome” (Hommes, 2002). While one can of course still tweek the model, and then examine the changes in outcomes, the approach is not accommodating to gaining an understanding of how these changes come about.

The economic dynamics literature takes a different route. Here, the objective is to build agent-based models that make it possible to study the joint dynamics of expectations and aggregate outcomes analytically (for an overview, see e.g. Hommes (2006) and Hommes and Wagener (2009)). To achieve this analytical tractability, the realistic but complex expectation formation process is stripped to its bare essentials. The models “may be viewed as simple, stylized versions of the more complicated ‘artificial markets’ and computationally oriented agent-based simulation models” (Hommes, 2006). They aim to accommodate the key distinctive features of realistic expectations like econometric models aim to accommodate the key stylized facts of empirical data. To ensure that the model has every chance of providing an accurate account of the mechanisms underlying empirical data, one may impose the condition that the expectations held by agents in the model are consistent with the price data generated by the model (i.e. that obvious tests for misspecification will not reject the models), see e.g. Hommes and Sorger (1998), Hommes and Rosser (2001)

and Sogner and Mitlochner (2002). “Economists once thought that behavior was either rational or impossible to formalize. We now know that models of bounded rationality are both possible and also much more accurate descriptions of behavior than purely rational models” (Barberis and Thaler, 2002).

1.4.2 Building econometric models

Here the objective is to best fit the data with a model that is as simple as possible, and with parameters that can be identified from the data. The econometric model cares not about the underlying economic or behavioral mechanisms, as long as the model is capable of describing the observed data. This means that the model and its parameters will typically not have an economic or behavioral interpretation attached. Stylized facts often function as the blue-print on which econometric models are build. If a new stylized fact is discovered, and deemed important, the model will be extended in a way that requires minimum changes yet provides a good fit.

For multivariate volatility models the number of parameters can easily rise very rapidly. The ‘holy grail’ is to find the model that is general, does not overfit the data, and yet is feasible in terms of estimation.

1.5 Overview of chapters

The thesis consists of five chapters in addition to this introductory chapter. The first part is concerned with fitting empirical data from financial markets, while the second part is concerned with rationalizing the empirical regularities of this data. Chapters 2 and 3 put forward an econometric model for (multivariate) financial volatility. Chapter 4 functions as a bridge between the two parts. It infers market expectations about financial volatility from empirical option price data. The last two chapters build a structural economic model. Chapter 5 develops an analytic framework designed to model the co-evolution of market expectations and prices. Chapter 6 considers some special cases by zooming in on a number of behavioral features.

Chapter 2 puts forward a new multivariate volatility model, which we will refer to as Generalized Orthogonal GARCH (GO-GARCH). It is a member of the ARCH-family, where ARCH stands for autoregressive conditional heteroskedasticity. The ARCH model arguably denotes the most popular specification to date for modelling the time-variation of financial volatility. It was introduced by Engle (1982) and generalized into GARCH

by Bollerslev (1986). The multivariate extension, multivariate GARCH, provides a model specification for the time-variation of the covariance matrix (which has time-varying variances on- and covariances off the diagonal). Estimation of multivariate GARCH models can be problematic as the number of unknown parameters involved tends to rise rapidly with the dimension.

The first general multivariate GARCH models put forward in the literature adopt a relatively large number of parameters which leads to convergence difficulties of estimation algorithms (see e.g. Bauwens *et al.*, 2006). New specifications are often determined by means of practical considerations. The challenge is to find a parameterization of the covariance matrix that is feasible in terms of estimation at a minimum loss of generality.

The GO-GARCH model denotes a natural generalization of the O-GARCH model, and is nested as a special case in the more general BEKK model. It accomodates the key stylized facts of multivariate volatility while maintaining feasibility. Both artificial and empirical examples are included to illustrate the new choice of model. For the published version of the chapter, see van der Weide (2002).

In Chapter 3 we propose a new estimation method for the factor loading matrix in GO-GARCH models. The method is based on eigenvectors of suitably defined sample autocorrelation matrices of squares and cross-products of returns. The method is numerically more attractive than likelihood-based estimation. Furthermore, the new method does not require strict assumptions on the volatility models of the factors, and therefore is less sensitive to model misspecification. We provide conditions for consistency of the estimator, and study its efficiency relative to maximum likelihood estimation using Monte Carlo simulations. The method is applied to European sector returns. For the published version see Boswijk and van der Weide (2011).

Embedded in option prices are market expectations regarding future volatility. While the assumption of rational expectations has been a popular paradigm, it is difficult to ignore the subjective nature of expectations. The objective of Chapter 4 is to make market expectations visible as they evolve over time, and to price options in line with prevailing expectations, be they rational or non-rational. We put forward an analytically convenient option pricing framework that accommodates both stochastic volatility and asymmetric volatility. Daily estimates of the implied pdf of volatility are obtained by estimating the option pricing model one day at a time. We do not impose too much structure on how expectations are updated over time, but allow market expectations to take their course. See Peters and van der Weide (2011) for the working paper version.

With Chapter 5 we move to the modelling of the expectation formation process. We propose a new analytic framework for studying the joint time-variation of expectations and prices. Beliefs distributions are defined on a beliefs space representing a continuum of possible strategies agents can choose from. Agents base their choices on past performances and re-evaluate strategies as new information becomes available. By considering individual choices as random variables, which is natural in a random utility framework, heterogeneity in beliefs can be seen to act as a ‘natural source of randomness’. Our framework gives rise to a random dynamical system (RDS), the stochastic properties of which are directly related to the time-varying beliefs distribution. We consider some asset pricing examples and discuss several conditions, that involve dependence among agents and unequal market impact, under which the randomness persists even as the number of agents tends to infinity. The chapter has been published as a working paper, see Diks and van der Weide (2003).

Chapter 6 considers a simple example of the framework introduced in Chapter four with the objective to investigate the effects on price dynamics of several behavioral assumptions: (i) herd behaviour; (ii) a-synchronous updating of beliefs; and (iii) heterogeneity in time horizons (memory) among agents. The benchmark model with many traders yields a random walk driven by news. Introducing herding is shown to modify the random walk to an ARIMA(0, 1, 1) process, which is observationally equivalent to a reduction of the number of market participants. In terms of returns the model predicts MA(1) structure with a negative coefficient. Asynchronous updating leads to an MA(1) model for returns with GARCH(1, 1) innovations, and predicts a relation between the ARCH and GARCH coefficients. Heterogeneity in memory leads to long-range dependence in returns. In the empirical section we perform a modest ‘reality check’ concerning the predicted sign of the MA coefficient and the relation between the ARCH and GARCH coefficients for exchange rate data. For the published version, see Diks and van der Weide (2005).