

MACROECONOMIC IMPLICATIONS OF BEHAVIOURAL FINANCE THEORIES

By

Rhys ap Gwilym

*A Thesis Submitted in Fulfilment of the Requirements for the Degree of Doctor
of Philosophy at Cardiff University*

**Economics Section
Cardiff Business School
Cardiff University**

2009

UMI Number: U584384

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U584384

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

DECLARATION

This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

Signed *[Signature]* Date 22/12/09

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of PhD.

Signed *[Signature]* Date 22/12/09

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated.

Other sources are acknowledged by footnotes giving explicit references.

Signed *[Signature]* Date 22/12/09

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed *[Signature]* Date 22/12/09

ABSTRACT

MACROECONOMIC IMPLICATIONS OF BEHAVIOURAL FINANCE

THEORIES

Rhys ap Gwilym

In this thesis I consider the extent to which macroeconomic theory and policy evaluation should be based upon behavioural models of human decision making.

I review the literature on decision making, and contrast it with the rational paradigm on which economic modelling is traditionally predicated. I also review that part of the macroeconomic literature which is based, explicitly or implicitly, on behavioural theories.

I develop a model of behavioural decision making in which investors base their portfolio decision on a choice between two simple heuristical forecasting rules. By simulating the model, I conclude that it can account for the observed history of the FTSE All-Share Index. By comparing this result with the ability of rational expectations models to account for historical asset prices, I conclude that behavioural theories of decision making do have a useful role in explaining macroeconomic time series.

Given that sub-rationality is important in helping to explain the macroeconomy, the question then arises as to whether there is scope for policy to correct for the misallocation of resources that is caused by this irrationality. I introduce the heuristical decision making model into a wider dynamic stochastic general equilibrium model of the entire economy. This allows me to assess whether using monetary policy to target asset price misalignments can enhance welfare. I find that in my particular model, a counter-intuitive ‘running with the wind’ monetary policy could enhance welfare. This result is clearly specific to the specification of decision making that I use, and runs counter to other intuitive arguments in favour of a ‘leaning against the wind’ policy. I, therefore, conclude that a systematic monetary policy response to asset mis-pricing is unlikely to enhance welfare.

Acknowledgements

I would like to thank Patrick Minford, my PhD supervisor, for all of his guidance and support during the research for this thesis. His ability to analyse my work and provide useful insight and concise critique has been of immense value.

I would also like to express my thanks to the wider economics community in Cardiff. In particular, I would like to thank my contemporaries on the PhD scheme and the scheme's coordinator, Max Gillman, for providing a stimulating environment in which to work.

I am also grateful to various delegates for feedback on the work I have presented at the Cardiff Economics workshops and at the Welsh Economics colloquium.

Contents

ABSTRACT	iii
Acknowledgements	iv
List of Tables	viii
List of Figures	ix
Chapter 1. Introduction	1
Chapter 2. Economic Decision Making	5
2.1. Rationality	8
2.2. Bounded Rationality, Heuristics and Biases	31
2.3. Rationality and Irrationality in Economic Modelling	59
2.4. Behavioural Considerations in Macroeconomics	81
Chapter 3. A Behavioural Explanation for Asset Prices	89
3.1. A Behavioural Asset Pricing Model	96

3.2. Basic Analysis of the Model	102
3.3. Disciplining the Model	112
3.4. The Data	115
3.5. Testing the Model	120
3.6. Conclusions	127
 Chapter 4. The Monetary Policy Implications of Behavioural Asset Bubbles	 131
4.1. Literature Review	133
4.2. The Model	143
4.3. Solving the Model	162
4.4. Preliminary Analysis of the Dynare sub-model	171
4.5. Results	178
4.6. Conclusions	190
 Chapter 5. Overall Conclusions	 193
 References	 200
 Appendix A. A Behavioural Asset Price Model	 209
A.1. Fortran code	209
A.2. Stochastic simulations	223
A.3. Results of the Stochastic Simulations	229

Appendix B. A DSGE Model with Asset Bubbles	232
B.1. Parameters for the Benchmark Model	232
B.2. The Dynare Sub-model: Equation listing	233
B.3. The Dynare Code	237
B.4. The Behavioural Sub-model: Equation listing	243
B.5. The MATLAB code	244
B.6. Robustness Tests	251

List of Tables

2.1 A solution to the St Petersburg paradox	18
2.2 Savage's analysis of the Allais paradox	25
3.1 Moments and time series properties of the simulated and actual FTSE series	122
4.1 Utility cost of various Taylor rules in the benchmark model	180
B.1 Utility cost of various Taylor rules in the model with no switching between forecasting rules	251
B.2 Utility cost of various Taylor rules in the model with rapid switching between forecasting rules	251
B.3 Utility cost of various Taylor rules in the model with a 1 percent uncertainty bound around the fundamentalist forecast	252
B.4 Utility cost of various Taylor rules in the model with a 10 percent uncertainty bound around the fundamentalist forecast	252

List of Figures

2.1 A basic utility function	17
2.2 Preferences under SEU	21
2.3 The Allais paradox	24
2.4 The ‘common consequence effect’	26
2.5 Kahneman and Tversky’s value function	44
2.6 Kahneman and Tversky’s weighting function	46
2.7 The sensory continuum	50
2.8 Deviations in Royal Dutch / Shell equity values	73
2.9 Prospective value of holding pure stock and pure bond portfolios	76
2.10 Prospective value of holding mixed portfolios	77
3.1 An example stochastic simulation	104
3.2 Sensitivity analysis in the deterministic model	109
3.3 Real FTSE 1963-2006	115
3.4 FTSE real quarterly rates of return and estimated population moments	116
3.5 Real UK Profits 1963-2006	118

3.6 Rate of change of real quarterly profits and estimated population moments	119
3.7 p-values and likelihood values for each model	125
4.1 Bordo and Jeanne (2002 [23]) optimal monetary policy	138
4.2 Solving the model	164
4.3 Impulse responses to a 1% productivity shock (i.e. $\varepsilon_1 = 0.01$)	172
4.4 Impulse responses to a 1% shock to the current asset price (i.e. $\varepsilon_1^q = 0.01$)	173
4.5 Impulse responses to a 1% shock to the expected future asset price (i.e. $\varepsilon_1^{q^F} = 0.01$)	175
4.6 Impulse responses to a 1% shock to the Taylor rule (i.e. $\varepsilon_1^{MP} = 0.01$)	176
4.7 The productivity process and rational asset price in a simulation of the model	184
4.8 The asset price in the rational model and in the behavioural model with a passive monetary policy	185
4.9 The asset price in the behavioural model with a ‘leaning against the wind’ monetary policy	186
4.10 The asset price in the behavioural model with a ‘running with the wind’ monetary policy	188

CHAPTER 1

Introduction

The question at the heart of this thesis is to what extent the field of macroeconomics should incorporate the theory that is being developed within the contemporary behavioural finance literature.

The importance of the psychological underpinnings of economic decision making has been a recurrent theme in macroeconomic literature. From Keynes' General Theory [60] in 1936 to Akerlof's Nobel prize lecture [4] in 2001, the most eminent of economists have highlighted the importance of behavioural issues to macroeconomic study.

However, despite their recurrence in macroeconomic literature, behavioural considerations have never taken centre stage. There are a number of possible reasons for this.

One reason is the difficulty of axiomising non-rational behaviour. It is the ease with which rationality can be defined, and the tractability with which it can be applied, that has made it such a popular assumption in macroeconomic modelling.

Other conceptions of decision making have regularly been criticised as ad hoc and lacking rigorous foundations.

Another reason is that, even if irrationality is essential in describing individual behaviour, it may well prove unimportant at the aggregate level. This could be the case if the irrational aspects of behaviour were random and uncorrelated, so that they cancelled out on aggregate. It could also be the case if irrationality were confined to a portion of the population, and its effects could be arbitrated away by rational agents. In such a world one individual's irrationality would provide an opportunity for other individuals to profit by intervening in ways which promote efficiency.

Even where behavioural themes have been discussed in the macroeconomic literature, much of this has been of an ad hoc nature. It has rarely been directly informed by insights from research in psychology.

The main questions that I address in this thesis are, firstly, whether there exist behavioural models that can explain macroeconomic data in a more complete way than the rational paradigm. Secondly, if departures from rationality are an important explanatory factor, what implications does this have for macroeconomic policy?

I begin this thesis by considering the mass of research that identifies the ways in which people actually make decisions. This research in cognitive psychology provides the basis on which economists might build their models. We shall see that although there is no overwhelming consensus on the nature of decision making, there is general agreement that people are not truly rational in the traditional sense that their decision making is based on a fully informed optimisation. The debate, rather, is concerned with the nature and extent of the departures from rationality that people display.

I continue in chapter 2 to examine how different views on the nature of decision making have informed economic modelling. This includes a review of the behavioural finance literature, the area of economics which has been most informed by psychology. Although macroeconomics has been less directly linked with contemporary advances in psychology, many developments in the field have been informed by reflections on human behaviour. I conclude chapter 2 with a review of the main theories in macroeconomics that have been formed on the foundations of behavioural considerations.

In chapter 3, I develop a model of asset prices under behavioural assumptions. The rational expectations, efficient markets paradigm has thrown up many well known puzzles when it comes to asset price dynamics. Rational models struggle to explain the equity premium, the volatility of asset prices, the high kurtosis of

stock returns, and a number of other features. Many behavioural researchers have argued that these anomalies can be explained by sub-rational behaviour. I test whether a simple, heuristical model of asset price determination can account for all of the time series properties that are apparent in asset price data. As we shall see, I find that I cannot reject my behavioural model as being the basis for real world asset price dynamics.

If sub-rational decision making is an important characteristic of financial markets, then there may well be room for welfare enhancing intervention in those markets. Irrational behaviour represents an additional source of risk in financial markets (often referred to as noise trader risk) and if interventions can be designed that mitigate or eliminate this risk, then efficiency will be enhanced. In chapter 4, I develop a New Keynesian dynamic stochastic general equilibrium model. However, I populate the model with agents who base their decision making on the sub-optimal rules suggested in chapter 3. I use this model to assess whether monetary policy can be used to correct for behavioural biases. I find that the complex and discontinuous nature of these biases makes it unreasonable to expect that a systematic policy is able to moderate their effects.

In chapter 5, I provide a summary of my findings, and consider their implications for further research.

CHAPTER 2

Economic Decision Making

Economics is the study of the allocation of scarce resources. It examines how consumers and firms interact with each other to determine that allocation. Clearly, we cannot get very far in analysing these interactions (the market) if we have no underlying theory of the behaviour and motivation of the individuals involved. Decision making theory, therefore, is one of the foundations on which all economic theory must be based.

The challenge for macroeconomists is to explain the behaviour of entire economies. Essentially, it is the challenge of aggregating the decisions of all the agents in the economy. Decision making is thus at the heart of macroeconomics, albeit augmented by the complexities of aggregation.

Models of decision making have always attempted to fulfil two different roles – to describe and to prescribe the decision making process. Descriptive, or positive, models attempt to give us an insight into how human beings actually arrive at a

decision outcome, whilst normative models focus on prescribing a set of behaviours against which we can judge actual behaviour. Psychologists are naturally concerned with both sides of this duality – they want to be able to explain actual behaviour as fully as possible, whilst at the same time having a standard against which to compare that behaviour.

As economists, we require a theory of decision making which can be used as the basis for our economic modelling. As such, we require a theory which can be integrated into a larger model of behaviours and interactions, and therefore a theory which can be reduced to a simple axiomatisation is our ultimate goal. For this reason, economists have generally shown greater interest in normative models of behaviour which, by their nature as a prescriptive set of rules, can always be presented simply.

However, economists are overwhelmingly concerned with explaining the world as it is. This is so that they can prescribe relevant policy, or forecast future events. The world as it might or should be may provide an interesting and relevant comparison for the world as it is, but is not the central concern of economics. For this reason, positive models of decision making are likely to provide a more sound basis for economic modelling. The problem confronting economists, therefore, is

to model decision making in a tractable way without simplifying away many of the important aspects of how real people make decisions.

In this chapter, I will begin by considering the two main paradigms of decision making theory; rationality and the heuristics and biases literature. I will review the axiomatisation of these two competing theories and examine how they have been incorporated into economic modelling generally. I will conclude by considering the specific ways in which decision making theory, and deviations from rationality, have influenced the development of macroeconomics.

2.1. Rationality

Normative models of human decision making are synonymous with the concept of rationality. It is the measure against which we assess the appropriateness of human decisions. In this section, I will begin by exploring the concept of rationality before considering how it has been incorporated into economic modeling. I will then reflect on some of the most fundamental objections to the rational paradigm.

2.1.1. The Concept of Rationality

According to Aristotle¹, rationality is the crucial characteristic that distinguishes human beings from other animals. The Oxford Companion to Philosophy defines rationality as “a feature of cognitive agents that they exhibit when they adopt beliefs on the basis of appropriate reasons” [53].

We can contrast rationality with either non-rationality or irrationality. An agent is non-rational when it is incapable of reasoning; a rock, a tree or, in Aristotle’s opinion, a dog. On the other hand, an agent who is capable of being rational but who regularly violates the principles of rational assessment is referred to as irrational.

We can differentiate beliefs or judgments on a similar basis. Beliefs that are contrary to the dictates of reason are irrational. However, an entire class of beliefs

¹"we declare that the function of man is a certain form of life, and define that form of life as the exercise of the soul's faculties and activities in association with rational principle" Nicomachean Ethics Book I, Ch. 7 [8]

exists that can be neither rational nor irrational because the beliefs are matters of taste and are not an appropriate subject of reason. These are usually referred to as preferences, and can be considered as non-rational.

So far, we only have the vaguest of notions. If rationality lies simply in the ability to provide reason, cause or justification for our beliefs and our actions, then it seems difficult to label anything as irrational. It is rare that we can provide no explanation of our behaviour, but few would argue that this is enough to make our behaviour rational. What we require is some form of measure of reasonableness. As the definition in the Oxford Companion states, reasons must be 'appropriate'.

2.1.1.1. Reasonableness. Reason is often considered as the internal coherence of the beliefs that underpin our actions. Indeed, this idea is central to Savage's concept of subjective probability, which I discuss later in this chapter. However, the link that is sometimes made between consistency of action and coherence of belief is a fallacious one. If I am inconsistent in my choice of action when faced with identical situations, this may be because my beliefs are incoherent, but it may equally be that my beliefs remain unchanged and that I have a preference for variety. The fact that I wear a pair of shorts one day and a pair of long trousers on another equally warm day may be a sign of irrationality, that I choose my clothing randomly, or it may be that I enjoy the variety. Consistency seems to provide an incomplete account of what is reasonable.

Logic is often suggested as a gauge of reason, and indeed it does provide a benchmark for reasonable belief in some circumstances. It is clearly irrational to believe at once that the world is round, R , and that it is not round, $\neg R$. Logic tells us that $(R \wedge \neg R)$ is a contradiction. However, logic's realm is narrow. It is concerned primarily with the field of certainty and demonstrative proof, whilst the majority of human life takes place in the realm of the merely probable. Logic can certainly help us to define what is reasonable, but it cannot provide the whole story.

Once again, Aristotle's works provide us with a good focus for our consideration of reasonableness. In his *Nicomachean Ethics*, he provides us with a consequentialist account of what is reasonable. An action is reasonable, and therefore rational, only if the agent does it in order to promote his 'eudaimonia'. A belief is reasonable only to the extent that it causes the agent to act reasonably. Eudaimonia is commonly translated as 'happiness', but literally means 'having a good guardian spirit', and in Plato and Aristotle's works is often associated with virtue. The most appropriate translation is probably something like 'human flourishing'.

Such a measure of reasonableness appears to be effective in modern usage. It is usual for us to contrast rational actions with those which are carried out in the name of emotion, faith, authority, or by arbitrary choice. In each of these cases I act either without carrying out an appropriate assessment of the consequences or in spite of the results of such an assessment.

Similarly, we contrast rational beliefs with those which are arrived at through emotion, faith, authority, or by arbitrary choice. Rational beliefs may be arrived at by accumulating relevant evidence or by providing a formal proof.

As I have already argued, formal proofs only apply to a narrow field. Formal logic and mathematics provide the clearest examples of where demonstrative proofs give rise to rigorous rules for deciding whether a proposition should be believed.

Most of human life, however, belongs to the realm of the probable, and as such the theory of probability has emerged as a guide to the reasonableness of the inferences and decisions that we make. Crucially, this entails that agents act in accordance with Bayesian rules of inference when assessing uncertain outcomes.

2.1.1.2. Bayesian inference. Bayes' theorem is crucial in informing decision making because it allows prior probabilities to be updated repeatedly in the light of new evidence to provide posterior probabilities. This is the essence of Bayesian inference, that prior probabilities can be modified in the light of data or empirical evidence in accordance with Bayes' theorem to yield posterior probabilities, which may then be used as prior probabilities for further updating in the light of subsequent data.

Bayes' theorem is easily proved from Bayes' rule, which in turn follows immediately from the fourth of Kolmogorov's axioms of probability. Kolmogorov's fourth axiom states that the probability of a conjunction of two events is equal to the

probability of the first event conditional upon the second occurring, multiplied by the unconditional probability of the second event:

$$\Pr(H \wedge D) = \Pr(H | D) \cdot \Pr(D)$$

Bayes' rule follows immediately:

$$\Pr(H | D) = \frac{\Pr(H \wedge D)}{\Pr(D)}$$

and, hence, Bayes' theorem:

$$\Pr(H | D) = \Pr(H) \cdot \frac{\Pr(D | H)}{\Pr(D)}$$

In other words, the conditional probability of an event H given an event D can be expressed as the product of the prior probability and a likelihood ratio. More generally, Bayes' theorem can be expressed as:

$$\Pr(H_n | D) = \Pr(H_n) \cdot \frac{\Pr(D | H_n)}{\sum_i [\Pr(H_i) \cdot \Pr(D | H_i)]}$$

where H_n is one of a set H_i of mutually exclusive and exhaustive events.

2.1.1.3. A definition of rationality. It is appropriate to conclude here with as formal a definition of rationality as possible. As we have seen, rationality is best viewed as a consequentialist concept. A rational agent acts as he does in

order to achieve some outcome. In Aristotle's prescription, he seeks to optimise his eudaimonia. I shall leave the examination of the definition of eudaimonia to later, save to note that its modern conception as 'utility' is of crucial importance to our discussion, and that all such conceptions are of a subjective, self-regarding function. Furthermore, Bayes' theorem makes clear the importance of the updating of probability judgments in the light of new information.

In short, we can define rational man as a self-regarding optimiser who makes full use of available information.

2.1.2. Axiomising Rationality: Subjective Expected Utility Theory

2.1.2.1. Expected Value. The problem of inductive inference has been a prominent topic in philosophical discourse at least since Aristotle. However, a formal, mathematical investigation of the problem only emerged at the end of the seventeenth century, when Pascal and Fermat exchanged letters about gambling problems. The first attempt at a unified theory of probability appeared in the form of Jacob Bernoulli's *Ars Conjectandi* (1713 [21]). The initial yardstick for reasonableness in decision making was to choose the alternative, or prospect, that maximises expected value:

$$\textit{expected value of prospect} = \sum x_i \cdot \pi_i$$

where π_i is the probability and x_i the value of the i^{th} certain outcome occurring.

This “principle of mathematical expectation” is well illustrated by Pascal’s famous ‘wager’ (Pascal 1660/1961 [74]). If God exists, a believer will enjoy everlasting happiness in the afterlife. We must weigh this outcome against the short-lived worldly pleasures that may be gained by non-believers. Pascal reasoned that as long as there is some finite possibility of God existing, however small, the expected value of the believer’s payoff will outweigh the expected value of the non-believer’s payoff. Rational self-interest dictates that we should forego the certain but finite worldly pleasures for the uncertain but infinite heavenly pleasures.

The definition of reasonableness in terms of expected value was soon exposed as inadequate on the basis of its incompatibility with the intuition of the majority of people, educated or not. The most famous conflict is the St Petersburg paradox, first publicised by Nicholas Bernoulli. The St Petersburg game is played by tossing a fair coin as many times as is necessary until it comes up heads. The game yields a prize equal to $\mathcal{L}2^n$, where n is the total number of tosses made. In other words, if a head comes up on the first toss the prize is equal to $\mathcal{L}2^1 = \mathcal{L}2$; if it takes two tosses to get a head the prize is $\mathcal{L}2^2 = \mathcal{L}4$; if it takes three attempts the prize will be $\mathcal{L}8$; and so on. Clearly, the probability of a head appearing on the first toss is $\frac{1}{2}$; the probability of needing to make a second toss and then getting a head is $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$; the probability of getting a head on the third attempt is $\frac{1}{8}$; and so on. By multiplying the value of each possible outcome by its probability and summing over all the outcomes, we find that the expected value of playing the St Petersburg

game is infinite:

$$\begin{aligned}\sum x_i \cdot \pi_i &= \left(2 \times \frac{1}{2}\right) + \left(4 \times \frac{1}{4}\right) + \left(8 \times \frac{1}{8}\right) + \left(16 \times \frac{1}{16}\right) + \dots \\ &= \infty\end{aligned}$$

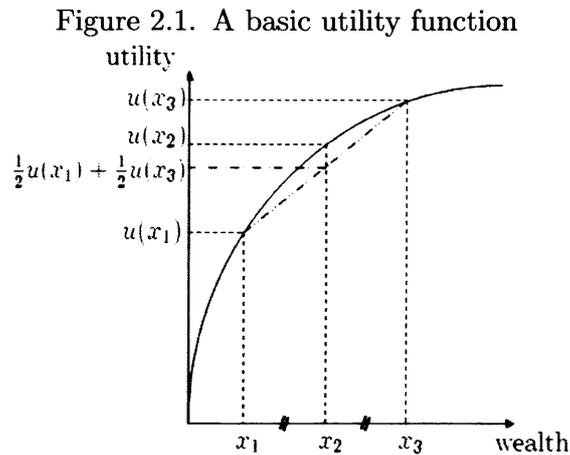
According to the “principle of mathematical expectation” it would be reasonable to pay any finite sum to play this game because the expected payout is infinite. Clearly, this result is counter-intuitive. Indeed, as Hacking (1980 [51]) claims, few of us would pay even £25 to enter such a game, let alone any finite sum.

2.1.2.2. The concept of utility. Daniel Bernoulli (1738/1954 [20]) introduced the concept of utility as an explanation of this divergence between rational theory and the prescription of apparent good sense. Bernoulli observed that the practice of calculating expected values implies that “no characteristic of the persons themselves ought to be taken into consideration; only those matters should be weighed carefully that pertain to the terms of the risk”. However, he recognised that whilst he would advise a rich man to pay 9,000 ducats to buy a lottery ticket which paid out 20,000 ducats with a probability of 50%, he would equally advise a pauper to sell the same lottery ticket for 9,000 ducats. He surmises that the reason for this is that there is a divergence between the monetary value of the lottery ticket and its subjective value, or utility, to each of the individuals: “the price of the item is

dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate”.

Bernoulli proposed changing the theory of decision making, suggesting that the yardstick for determining the reasonableness of a decision was in its propensity to maximise expected utility rather than expected value. Central to Bernoulli’s notion of utility is the concept of diminishing marginal utility. As can be seen from the utility function in figure 2.1, successive increases in wealth do not yield proportional increases in utility. As an individual’s wealth increases from x_1 to x_2 , their level of utility rises from $u(x_1)$ to $u(x_2)$. However, if the same individual’s wealth increases by the same amount again, this time from x_2 to x_3 , then their utility will only increase from $u(x_2)$ to $u(x_3)$, a much smaller amount than with the original increase. It is the concavity of the utility function that gives rise to this diminishing marginal utility of wealth. It is this characteristic of the utility function which also implies risk aversion. As can be seen clearly from figure 2.1, $u(x_2)$ is greater than $\frac{1}{2}u(x_1) + \frac{1}{2}u(x_3)$ even though x_2 is equal to $(\frac{1}{2}x_1 + \frac{1}{2}x_3)$. In other words, the individual would prefer the certainty of x_2 to a risky prospect which has the same expected value.

Bernoulli in fact proposed the logarithmic utility function that is still popular today, and which is not a dissimilar shape to that in figure 2.1. Given natural log utility, we can easily explain the St Petersburg paradox. Table 2.1 shows that, whilst the expected value of the St Petersburg payoff is infinite, the utility associated with the bigger prizes decays quickly and so the expected utility for



an individual with a natural log utility function would be 1.39. This is the same utility level as the individual will associate with the certainty of £4 ($\ln 4 = 1.39$). In other words, an individual with a natural log utility function would not pay more than £4 to play the St Petersburg game. This may well seem a reasonable stake for a risk averse individual, given that there is still a 50% chance of making a loss and only a 25% chance of making a profit.

It is important to note, however, that with a log utility function the St Petersburg paradox can be amended to produce a random act with infinite expected utility that, again, no one would really prefer to the status quo. Savage (1954 [82]) provides the following less elaborate but equally counter-intuitive example. A man with initial wealth of £1 million and whose utility is simply the natural log of his wealth would be indifferent about flipping a fair coin that would decide whether his wealth increased to £100 million or fell to £10,000. Even though log utility may

Table 2.1. A solution to the St Petersburg paradox

Number of tosses	Probability, π_i	Prize, x_i	Utility, $u(x_i) = \ln(x_i)$	$x_i \cdot \pi_i$	$u(x_i) \cdot \pi_i$
1	$\frac{1}{2}$	2	0.69	1	0.35
2	$\frac{1}{4}$	4	1.39	1	0.35
3	$\frac{1}{8}$	8	2.08	1	0.22
4	$\frac{1}{16}$	16	2.77	1	0.17
5	$\frac{1}{32}$	32	3.47	1	0.11
6	$\frac{1}{64}$	64	4.16	1	0.06
7	$\frac{1}{128}$	128	4.85	1	0.04
8	$\frac{1}{256}$	256	5.55	1	0.02
9	$\frac{1}{512}$	512	6.24	1	0.01
10	$\frac{1}{1024}$	1024	6.93	1	0.01
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
Expected Value = $\sum x_i \cdot \pi_i$				∞	
Expected Utility = $\sum u(x_i) \cdot \pi_i$					1.39

be a good approximation to the shape of a utility function for moderate ranges of wealth, a utility function must have an upper and a lower bound if it is to avoid criticisms of the type inherent in the St Petersburg paradox.

2.1.2.3. The formalisation of expected utility theory. Von Neumann and Morgenstern (1944 [95]) formalised and axiomatised the theory of expected utility. They showed that an individual whose preferences satisfy certain well-accepted

axioms would choose between different prospects as though maximising expected utility.

The first axiom of expected utility theory is the ordering axiom, which requires that the individual should have a preference ordering over the set of all conceivable prospects. Furthermore, that ordering should be complete, reflexive and transitive.

The second axiom, the continuity axiom, requires that for any three prospects p , q and r such that p is preferred to q and q is preferred to r ($p \succ q \succ r$), then there must be some compound prospect of p and r in some probabilities λ and $(1 - \lambda)$ such that it is indifferent to q .

Finally, the independence axiom states that the evaluation of a prospect is not affected if some element in it is replaced by another element with which it is indifferent. In choosing between prospects, agents decide on the basis of what distinguishes them and not on the basis of what they have in common.

If these three axioms hold then von Neumann and Morgenstern (1944 [95]) show that it is possible to assign to every prospect p a real-valued utility index $u(p)$ such that the utility index $u(\cdot)$ represents the individual's preference ordering. Furthermore, if prospect p yields consequences x_1, x_2, \dots, x_n with probabilities $\pi_1, \pi_2, \dots, \pi_n$, then $u(p) = \pi_1 \cdot u(x_1) + \pi_2 \cdot u(x_2) + \dots + \pi_n \cdot u(x_n)$.

For simplicity, I will consider the set of prospects where there are three pure consequences x_1, x_2, x_3 such that $x_1 \prec x_2 \prec x_3$. A typical prospect may be written as (π_1, π_2, π_3) where π_1, π_2 and π_3 are the probabilities of x_1, x_2 and x_3 in the prospect. Since $\pi_2 = (1 - \pi_1 - \pi_3)$ we may suppress π_2 and thus represent the

prospects in (π_1, π_3) space, as in figure 2.2. The utility index for any prospect (π_1, π_3) is given by:

$$u(\pi_1, \pi_3) = \pi_1 \cdot u(x_1) + (1 - \pi_1 - \pi_3) \cdot u(x_2) + \pi_3 \cdot u(x_3)$$

There must be a set of prospects which generate the same expected utility, say u^* . These prospects form an indifference curve in (π_1, π_3) space with equation:

$$u^* = \pi_1 \cdot u(x_1) + (1 - \pi_1 - \pi_3) \cdot u(x_2) + \pi_3 \cdot u(x_3)$$

By rearranging:

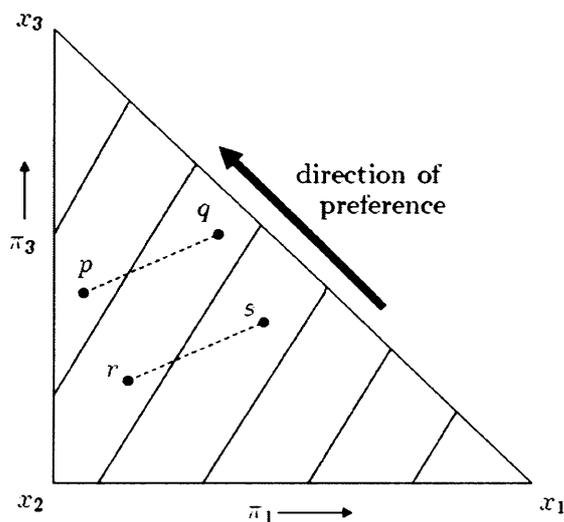
$$\pi_1 [u(x_1) - u(x_2)] + \pi_3 [u(x_3) - u(x_2)] = u^* - u(x_2)$$

If we differentiate π_3 with respect to π_1 , we get:

$$\frac{d\pi_3}{d\pi_1} = \frac{u(x_2) - u(x_1)}{u(x_3) - u(x_2)}$$

We can see that the gradient of the indifference curve is independent of u^* and, therefore, that under von Neumann and Morgenstern's axioms the family of indifference curves that an individual associates with any set of pure consequences must be linear and parallel. In the triangle in figure 2.2, showing the set of conceivable prospects in (π_1, π_3) space, the parallel lines are indifference curves with the arrow showing the direction of preference (pure consequence x_3 is preferable to pure consequence x_1 by assumption).

Figure 2.2. Preferences under SEU



2.1.2.4. Subjective probability. Savage (1954 [82]) provides the complete axiomatisation of subjective expected utility theory (SEU). He introduces the concept of subjective probability into von Neumann and Morgenstern's theory, which allows the theory to be applied to situations in which there is no objectively known probability. He shows that under certain axioms, preferences can be represented by the expectation of a utility function, this time weighted by the individual's subjective probability assessment.

2.1.3. Criticisms of the Rational Paradigm

2.1.3.1. The Allais paradox. The implications of von Neumann and Morgenstern's theory is that we can construct a utility function by observing preferences locally because if we know the shape of the indifference curves locally then we know the shape of all the indifference curves. Hence, if we know an individual's preference between p and q in figure 2.2, we know his preference between r and s must be the same because rs is parallel to pq . This implication of SEU theory was first highlighted by Maurice Allais (1953 [7]), and has become popularly known as the Allais paradox because it conflicts with empirical studies of people's preferences.

Allais (1953 [7]) illustrates his point with the following pair of questions:

(1) Would you prefer prospect A or prospect B?

Prospect A: the certainty of receiving 100 million francs

Prospect B: 10% chance of winning 500 million

89% chance of winning 100 million

1% chance of winning nothing

(2) Would you prefer prospect C or prospect D?

Prospect C: 11% chance of winning 100 million

89% chance of winning nothing

Prospect D: 10% chance of winning 500 million

90% chance of winning nothing

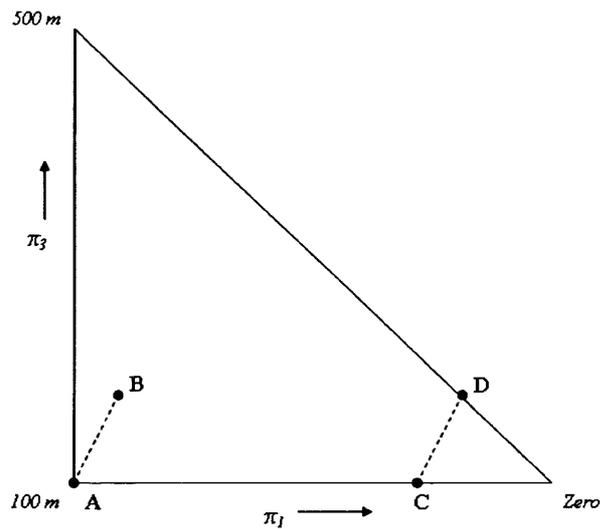
We can represent Allais' prospects in the same manner as figure 2.2, where x_1 is zero francs, x_2 is 100 million francs and x_3 is 500 million francs. These pure consequences adhere to our assumed preference ordering, $x_1 \prec x_2 \prec x_3$. We can then represent the prospects in (π_1, π_3) space as follows:

Prospect A:	(0.00,0.00)
Prospect B:	(0.01,0.10)
Prospect C:	(0.89,0.00)
Prospect D:	(0.90,0.10)

These prospects are shown in figure 2.3. Clearly, the lines AB and CD are parallel, and so according to SEU theory, if A is preferred to B then C must be preferred D and vice versa. In empirical studies, however, people generally prefer A to B, but prefer D to C. Kahneman and Tversky (1979 [56]) report the results of a variation on the above choice problem. In a trial involving 72 respondents, 82% expressed a preference for prospect A over prospect B whilst 83% expressed a preference for prospect D over prospect C. Clearly, such empirical results bring into question the theory expounded above.

Savage (1954 [82]) argues that examples such as the one cited above have great intuitive appeal, but that the standard reaction is simply an irrational response. He admits that his immediate response to Allais' problem was to express preference for prospects A and D. However, he argues that on closer inspection he is compelled to reverse his choice in respect of the second question, and that this amounts

Figure 2.3. The Allais paradox



to the correction of an error. He rephrases the same problem as a lottery with a hundred numbered tickets and with prizes as shown in table 2.2. If a ticket numbered between 12 and 100 is drawn then it does not matter, in either case, which prospect is chosen. Focusing, therefore, merely on the outcomes if one of the tickets numbered between 1 and 11 is drawn, it is apparent that the same decision needs to be made in either situation (1) or situation (2). The decision in both situations is whether to take an outright gift of 100 million francs or to take a ten to one on chance of winning 500 million francs. Savage concludes that he would prefer the sure thing to the gamble, and hence prefers A to B and C to D.

Savage's analysis here is a clear defence of von Neumann and Morgenstern's independence axiom. He considers lottery tickets 12 to 100 in isolation and; having

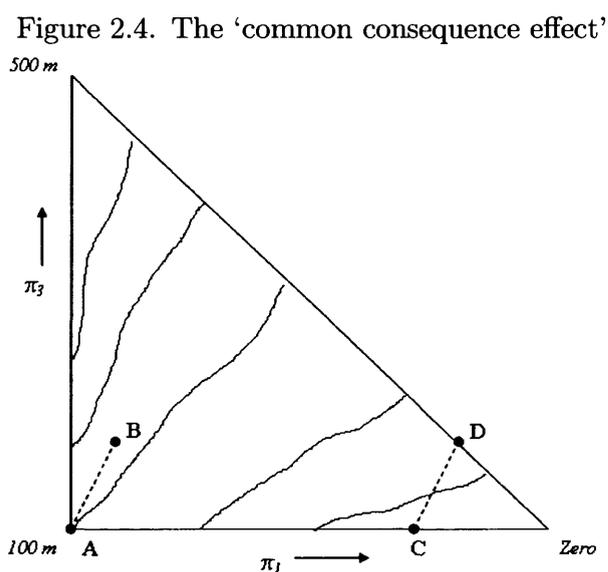
Table 2.2. Savage's analysis of the Allais paradox

		Ticket Number		
		1	2-11	12-100
(1)	Prospect A	100m	100m	100m
	Prospect B	0	500m	100m
(2)	Prospect C	100m	100m	0
	Prospect D	0	500m	0

disregarded them because they do not provide a basis for distinguishing between competing prospects; he then considers tickets 1 to 11 in isolation.

The independence axiom entails that if prospect p^* is preferred to prospect p , then the mixture $[\alpha p^* + (1 - \alpha)p^{**}]$ is preferred to $[\alpha p + (1 - \alpha)p^{**}]$ for all $\alpha > 0$ and p^{**} , and it is this that ensures the linearity of the indifference curves. This property can be interpreted as follows:

“In terms of the ultimate probabilities over the outcomes $\{x_1, x_2 \dots x_n\}$, choosing between the mixtures $[\alpha p^* + (1 - \alpha)p^{**}]$ and $[\alpha p + (1 - \alpha)p^{**}]$ is the same as being offered a coin with a probability of $(1 - \alpha)$ of landing tails, in which case you will obtain the lottery p^{**} , and being asked before the flip whether you would rather have p^* or p in the event of a head. Now either the coin will land tails, in which case your choice won't have mattered, or else it will land heads, in which case you are ‘in effect’ back to a choice between p^* or p , and it is only ‘rational’ to make the same choice as you would before” (Machina 1987 [66])



Clearly, the argument here is that there can be no complementarities between the different probability mixtures in the prospects.

Machina (1987 [66]) argues that the empirical evidence is not easy to dismiss and that we need to reassess the theory rather than reassess our initial reaction to Allais' problem. He argues that, just as with goods, there can be complementarities in the probability mixtures of prospects. The 'common consequence effect' entails that the better off an individual would be from winning 100 million francs in prospect A, the more risk averse they become over the gamble in prospect B. In other words, as one moves up and left in the triangle, in figure 2.3, towards one's preferred point, risk aversion increases, and the indifference curves become steeper.

By dropping the independence axiom, we can get non-linearity in the indifference curves, as in figure 2.4, and this would explain the Allais paradox.

Quite why we might have complementarities in the probability mixtures of prospects remains to be explained. I will leave this to the discussion of the heuristics and biases paradigm below.

2.1.3.2. Rabin. Matthew Rabin (2000 [75]) provides a more contemporary critique of subjective expected utility theory. He shows that if we posit a utility function which is strictly increasing and weakly concave across all values of wealth, then observed levels of risk aversion over bets with small stakes imply unrealistically large risk aversion over bets with large stakes. As an example, he shows that if we know that a risk-averse person turns down an evens bet of losing \$100 or gaining \$105, then SEU theory implies that the same individual will turn down an evens bet of losing \$4,000 or gaining \$635,000. This is clearly incompatible with observed behaviour. The intuition for this property of SEU theory is that turning down a small stakes bet with positive expected value means that the marginal utility of money must diminish very quickly for small changes in wealth. When this diminution is iterated over large changes in wealth, it gives rise to an unrealistically high rate at which the value of money deteriorates.

2.1.3.3. The Ellsberg paradox. Our discussion so far has focused on anomalies between observed behaviour and the normative SEU model in situations where the possible outcomes are well defined. The alternative prospects which Allais offered his subjects all had well specified probability distributions of outcomes. The Ellsberg paradox draws a distinction between circumstances in which the probability distribution of outcomes is known (risky outcomes) and those circumstances in which the distribution is poorly defined or unknown (uncertain outcomes).

Uncertainty can be thought of as a second dimension of risk. If we are asked to draw a card from what we know to be a conventional pack of playing cards, we do not know what colour the card will be but we do know with certainty that the probability of red is $\frac{1}{2}$ and the probability of black is $\frac{1}{2}$. However, if we are asked to draw a card from a bunch of playing cards which someone else has put together, we do not know how many of each colour is in the bunch and so do not know the exact distribution of possible outcomes. In this case we must form a subjective distribution of possible outcomes which will depend on our beliefs about various parameters, such as the tastes of the individual who selected the bunch. According to Ellsberg (1961 [34]), we are not only averse to the risk inherent in choosing a card from a distribution, but we are also averse to the extra dimension of risk inherent when we do not know for certain the distribution of possible outcomes. Barberis and Thaler (2002 [14]) refer to this as ambiguity aversion.

Ellsberg (1961 [34]) describes the following experiment. Suppose that there are two urns. Urn 2 contains a total of 100 balls, 50 red and 50 blue. Urn 1 also

contains 100 balls, again a mix of red and blue, but the subject does not know the proportion of each. Subjects are then asked to choose one of the following two prospects, each of which involves a possible payment of £100, depending on the colour of a ball drawn at random from the relevant urn:

Prospect A: a ball is drawn from Urn 1, you win £100 if red, nothing if blue

Prospect B: a ball is drawn from Urn 2, you win £100 if red, nothing if blue

Subjects are then asked to choose between following two prospects:

Prospect C: a ball is drawn from Urn 1, you win £100 if blue, nothing if red

Prospect D: a ball is drawn from Urn 2, you win £100 if blue, nothing if red

Prospect B is typically preferred to A, while D is preferred to C. These choices are inconsistent with SEU. The choice of B implies a subjective probability that fewer than 50% of the balls in Urn 1 are red, while the choice of D implies the opposite. The experiment suggests that people dislike uncertain outcomes more than they dislike straightforward risky outcomes where the distribution of possible outcomes is known.

2.1.3.4. Summary. Subjective expected utility theory is the seminal axiomatisation of rationality. It provides a simple framework on which to base decision making. The Allais and Ellsberg paradoxes and Rabin's critique provide serious objections to the use of SEU theory as a positive theory of human decision making.

However, there is a general consensus that it still provides the most appropriate normative measure against which decisions may be assessed.

2.2. Bounded Rationality, Heuristics and Biases

Allais and Ellsberg's criticisms of SEU are based on perceived weaknesses in the axiomatisation of that theory and its resultant failure to describe people's preferences for risky prospects. Here, we consider a different type of objection, one that is aimed at the root of rational theory. The theories of bounded rationality and of the heuristics and biases paradigm do not merely question the specification of SEU; they deny the idea that decision making is aimed at any form of optimisation.

2.2.1. Bounded rationality

The phrase 'bounded rationality' was first introduced by Herbert A. Simon in the mid-1950s, and the theory was initially associated with the field of artificial intelligence. Simon (1956 [85]) notes that "however adaptive the behaviour of organisms in learning and choice situations, this adaptiveness falls far short of the ideal of 'maximisation' postulated in economic theory. Organisms adapt well enough to 'satisfice'; they do not, in general, 'optimise'".

Models of bounded rationality take into account the limitations of knowledge and computational capacity in human decision making. They take into account how actual decision making processes influence the decisions that are reached. They do not reject the concept of rationality entirely in that they continue to view agents as having goals which they are motivated to fulfil. However, the

optimising behaviour associated with full rationality is rejected as too expensive computationally and in terms of the resources of knowledge required.

Simon (1987 [86]) contrasts boundedly rational models with subjective expected utility theory (SEU) in three ways. Firstly, SEU theory holds that choices are made from amongst a given set of alternatives. Boundedly rational models, on the other hand, recognise that in reality the search for different possible alternatives is costly. Therefore, these models must postulate processes for generating these alternatives. Secondly, SEU theory assumes that each of the possible alternatives has a subjectively known probability distribution of outcomes. Theories of bounded rationality deny the innate knowledge of these distributions, and so they must introduce procedures for estimating them or strategies for dealing with uncertainty that do not assume knowledge of probabilities. Finally, SEU theory holds that agents make choices in such a way as to maximise the expected value of their utility function, which is given a priori. Bounded rationality, however, holds that optimisation is too demanding of knowledge and computational capacity to be an effective tool for choice. Hence, theories of bounded rationality must postulate an alternative choice mechanism and this often involves a satisficing strategy.

2.2.2. Kahneman and Tversky's Heuristics and Biases

A more concerted challenge to the idea that rationality provided an effective model of human decision making was established by a series of papers published in the

early 1970s by Daniel Kahneman and Amos Tversky. They argued that rather than following rational rules, people use straightforward rules of thumb, heuristics, in assessing frequencies and probabilities. Furthermore, they argued that these heuristics lead to systematic biases in decision making when compared to the normative (rational) model.

Their seminal paper (Tversky and Kahneman 1974 [91]) draws attention to three heuristics which they claim can explain human decision making under uncertainty. These heuristics, and the biases to which they give rise, are discussed below.

2.2.2.1. Representativeness. People order probability and similarity in exactly the same way. In assessing the probability that an individual belongs to a particular group, people assess the degree to which the individual is representative (or similar to the stereotype) of that group.

Kahneman and Tversky illustrate this effect with the use of a description of an individual: "Steve is very shy and withdrawn, invariably helpful, but with little interest in people, or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail". According to the representativeness heuristic, people will assess the probability of Steve belonging to a particular occupational group from among a list of possibilities according to

how similar Steve is to their stereotypical image of the relevant occupation. This leads to a number of biases:

- (1) Insensitivity to prior probability of outcome. The base rate frequency of outcomes has no effect on representativeness but does affect probability. In the above example, people will judge Steve as more likely to be a librarian than a salesman because his description is more similar to the stereotype of a librarian, even though the number of salesmen in the population in general far exceeds the number of librarians. Similarity is not contingent upon prior probability, but posterior probability is.
- (2) Insensitivity to sample size. People assess the probability of a sample result by the similarity of this result to the corresponding population parameter, and hence ignore the size of the sample. Kahneman and Tversky (1972 [54]) report the results of an experiment where subjects were asked to assess the probability of a random sample of men having an average height of over six foot. Subjects assigned the same value to samples of 1000, 100 and 10 men, even though probability theory tells us that this is incorrect.
- (3) Misconceptions of chance. People expect that a sequence of events generated by a random process will represent the essential characteristics of that process even when the sequence is short. For example, when a coin is tossed six times people expect the outcome to represent the unbiased nature of the coin. Hence, people regard $H - T - H - T - T - H$ as

a more likely sequence than $H - H - H - T - T - T$, which does not 'look random' and so does not represent the fairness of the coin as well as the first sequence (Kahneman and Tversky 1972 [54]).

- (4) Insensitivity to predictability. When people make predictions about a future outcome based on a current description, they will base the prediction on the representativeness of the description regardless of the predictability of the future event. Hence, predictions about a remote future criterion are as variable as the evaluation of the current information. Kahneman and Tversky (1973 [55]) describe a study in which subjects were presented with several descriptions of student teachers giving practice lessons. They were then asked to firstly evaluate the quality of a lesson and secondly to predict the standing of the student teacher five years hence. The judgments made in each case were almost identical, even though probability theory would suggest that performance in one class now should not be a very good predictor of performance five years ahead.
- (5) The illusion of validity. The internal consistency of a pattern of inputs is a major determinant of one's confidence in predictions based on these inputs. However, consistent patterns are most often observed when the inputs are highly correlated, and therefore redundant.
- (6) Misconceptions of regression. People do not develop correct intuitions about regression towards the mean. It is incompatible with the belief

that the predicted outcome should be maximally representative of the input, and therefore just as extreme.

2.2.2.2. Availability. People assess the frequency of a class or the probability of an event by the ease with which instances or occurrences can be brought to mind.

This leads to a number of biases:

- (1) Bias due to the retrievability of instances. A class whose instances are easily retrieved will appear more numerous than a class of equal frequency whose instances are less easily retrieved. For example, Tversky and Kahneman (1973 [89]) report an experiment in which subjects heard a list of people's names and were then asked to judge whether the list contained more male or female names. In some lists, the men were more famous than the women and in other lists this was reversed. In each case, the subjects wrongly judged that the sex which had the more famous names was the more numerous.
- (2) Bias due to the effectiveness of a search set. Different tasks elicit different search sets and some searches are much easier than others. For example, because it is easier to search for words which have the letter r at the beginning of the word (e.g. road) than to search for words which have the letter r as the third letter (e.g. cart), most people erroneously judge that

there are more words of the first type than of the second type (Tversky and Kahneman 1973 [89]).

- (3) Biases of imaginability. Some outcomes are more easily imagined than others. For example, it is much easier to imagine two people from a group of ten than it is to imagine eight people from the group of ten. Of course, it is possible to combine two from ten in exactly the same number of ways as it is to combine eight from ten (each permutation of two leaves you with a unique combination of people left over). However, Tversky and Kahneman (1973 [89]) report that when subjects were asked to estimate the number of permutations, the median estimate of the number of combinations of two people was 70, whilst that for eight people was 20 (the correct answer in each case is 45). In the same way it is easy to explain people's aversion to 'adventurous activities', because the ease with which disasters are imagined need not reflect their actual likelihood.
- (4) Illusory correlation. The judgment of the frequency with which two events co-occur may be biased by an associative bond between them.

2.2.2.3. Adjustment and anchoring. People make estimates by starting from an initial value and then adjusting. Different starting points, possibly suggested by the formulation of the problem, might therefore lead to different estimates. The following biases are associated with the adjustment and anchoring heuristic:

- (1) Insufficient adjustment. Adjustments are typically insufficient. Two groups of school pupils were asked to estimate a numerical expression within five seconds. One group was asked to estimate: $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$. Another group estimated: $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$. The median estimate for the first group was 2,250 and that for the second group was 512. The correct answer is 40,320 (Tversky and Kahneman 1974 [91]).
- (2) Biases in the evaluation of conjunctive and disjunctive events. People tend to overestimate the probability of conjunctive events and to underestimate the probability of disjunctive events. The probability of the elementary event provides a natural starting point for the estimation of either and, since adjustment is typically insufficient, the final estimates remain too close to the probabilities of the elementary events in both cases.

2.2.3. Axiomising Heuristics: Prospect Theory

Kahneman and Tversky's heuristics literature provides an interesting insight into how human beings make judgments and, consequently, decisions. However, it provides a purely descriptive account of how individuals do so. It does suggest common biases which afflict all of us, but these are generally vague and ill-defined.

Prospect theory represents Kahneman and Tversky's (1979 [56]) attempt to provide a formal, axiomised model of human decision making². Prospect theory is still unequivocally a descriptive model, but it attempts to better define the commonality in the biases that are predicted by their previous work, and provide a unified theory of decision making. Their criticisms of SEU theory are all based on empirical studies and, hence, it is the behaviour which they have observed during these studies that they are focused on explaining. Providing a normative model which prescribes how decisions should be made is not their purpose here:

“in order to accommodate the effects [observed], we are compelled to assume that values are attached to changes rather than to final states, and that decision weights do not coincide with stated probabilities. These departures from expected utility theory must lead to normatively unacceptable consequences.” (Kahneman and Tversky 1979 [56])

They distinguish two distinct stages in the decision making process. The first is an editing phase which reformulates the options so as to simplify the subsequent evaluation phase.

2.2.3.1. The editing phase. The editing phase consists in the application of the following operations:

²Some of the core ideas inherent in prospect theory (notably, the existence of a reference point) had been presented much earlier by Markowitz (1952 [67]). However, prospect theory has come to overshadow this work, and is now synonymous with the idea of ‘non-expected utility theory’.

- (1) Coding. This essentially consists in locating a reference point against which gains and losses can be defined. Usually this will correspond with the individual's current asset position, but it may be affected by the way in which offered prospects are framed. For example, where one of the prospects involves a certain gain, this may be amalgamated with current assets to determine the reference point.
- (2) Combination. Prospects may be simplified by combining the probabilities of identical outcomes. For example, if a prospect involves winning £100 if a coin lands heads within two tosses, the 0.5 probability of winning on the first toss and 0.25 probability of winning on the second toss can be combined to give a simpler prospect of a 0.75 probability of winning £100.
- (3) Segregation. Prospects may be simplified by segregating a riskless component from the risky component. For example, a prospect where there is an 80% chance of winning £300 and a 20% chance of winning £200 may be separated into a sure gain of £200 and an 80% chance of winning £100.
- (4) Cancellation. When the offered prospects have components in common, these may be discarded. For example, a choice between the following prospects:

Prospect A: win £100 if a die lands 1,2 or 3; lose £100 if 4 or 5; win £200 if 6

Prospect B: win £100 if a die lands 1,2 or 3; nothing if 4 or 5; win £50 if 6

can be simplified by ignoring the common outcome-probability pairs to:

Prospect A: 33.3% chance of losing £100; 16.7% chance of winning £200

Prospect B: 16.7% chance of winning £50

5. Simplification. Probabilities or outcomes may be rounded to simplify evaluation. For example the prospect of winning £101 with a probability of 0.49 may be treated as the prospect of winning £100 with probability 0.50.
6. Detection of dominance. The available prospects are checked to find any which are strictly dominated by another and they are then dropped without further evaluation.

Some of the editing operations conflict with others and, hence, the same set of original prospects may result in varying sets of edited prospects. For example, in the case of the following prospects:

Prospect A: 20% chance of winning £100; 49% chance of winning £199

Prospect B: 15% chance of winning £100; 51% chance of winning £201

prospect A will appear to dominate prospect B if the second constituent of each is simplified to a 50% chance of winning £200. The order in which operations are applied will be crucial in determining the final set of edited prospects. Therefore, the context in which the original prospects are offered may affect the final choice

by influencing the ordering of operations. This is the all important concept of framing.

2.2.3.2. The evaluation phase. After the editing phase, the decision maker moves on to evaluate the simplified set of prospects, and to choose the one of greatest value (Kahneman and Tversky talk of the ‘value’ of an edited prospect, and derive a ‘value’ function. I will remain consistent with their semantics in this section. However, it is important to recognise that Kahneman and Tversky’s ‘value’ function is very different to the traditional notion of expected (monetary) value).

The expected (monetary) value of a prospect is simply the sum of the products of the possible outcomes and their probabilities:

$$EV = \sum x_i \cdot \pi_i$$

SEU theory defines expected utility as being non-linear in the values of each outcome:

$$EU = \sum u(x_i) \cdot \pi_i$$

Prospect theory treats the overall value of a prospect as non-linear not only in the values of each outcome, but also as non-linear in the associated probabilities. The value of a prospect is the sum of the products of functions of the possible outcomes

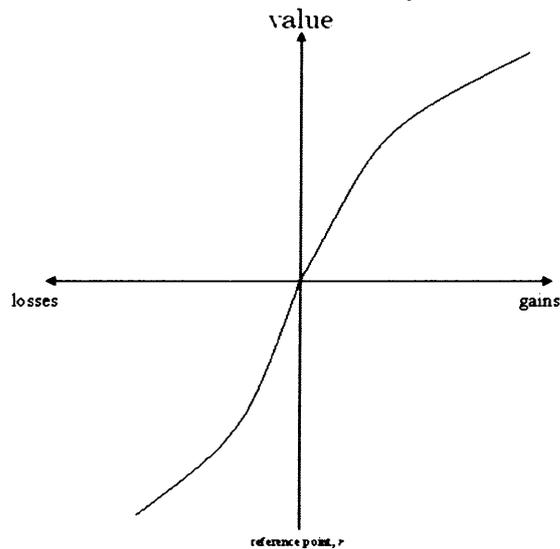
and their probabilities:

$$V = \sum v(x_i) \cdot w(\pi_i)$$

The value function: The $v(\cdot)$ function associates with each monetary value, x_i , a subjective value, $v(x_i)$. This is similar to the utility function in SEU theory, but differs in two significant ways. The first is that the monetary value, x_i , is generally regarded in SEU theory as the final asset position of the individual. In prospect theory, outcomes are defined relative to a reference point, r , determined during the editing phase. Therefore, in order to remain consistent in our terminology, I should really define the value function as a function of changes in wealth relative to the reference point, $v(x_i - r)$.

The second difference is that Kahneman and Tversky prescribe a specific shape to the value function. They propose a concave function for gains, as with SEU theory. For losses, however, they propose that the value function is convex. To explain this they draw an analogy with discriminating between temperatures. Whilst it is relatively easy to discriminate between a change of 3° and a change of 6° in room temperature, it is much more difficult to discriminate between a change of 13° and a change of 16°. This is true regardless of whether the change in question is an increase or a decrease. If we apply this same principle to the assessment of monetary values, then the difference between a gain of £100 and a gain of £200 will

Figure 2.5. Kahneman and Tversky's value function



appear greater than the difference between a gain of £1100 and a gain of £1200. This clearly implies the risk averse behaviour that expected utility theorists have always prescribed. However, if the difference between a loss of £100 and a loss of £200 similarly appears greater than the difference between a loss of £1100 and a loss of £1200, then this implies risk seeking behaviour where the choice is between two losses. Such a function, as illustrated in figure 2.5, is consistent with Kahneman and Tversky's empirical evidence.

The function is shown as steeper for losses than for gains because, again, empirical evidence suggests that losses loom larger than gains. People find symmetric bets (for example a 50% chance of winning £10 and 50% chance of losing £10) unattractive, and they find them more unattractive the greater the stake.

The weighting function: The $w(\cdot)$ function associates with each probability, π_i , a decision weight, $w(\pi_i)$, which reflects the effect of the given probability on the overall value of the prospect. $w(\pi_i)$ is an increasing function of π_i with $w(0) = 0$ and $w(1) = 1$ (i.e. impossible outcomes are ignored and the function is normalised so that $w(\pi_i)$ is the ratio of the weight associated with the probability π_i and the weight associated with certainty).

Based, again, on empirical observation, Kahneman and Tversky conclude that very low probabilities are overweighted. This would help explain why people might prefer a lottery ticket with very poor odds of paying out to the expected value of the ticket, and why people will buy fair insurance for events with a tiny probability of happening.

Although $w(\pi_i) > \pi_i$ for low probabilities, for all $0 < \pi_i < 1$, $[w(\pi_i) + w(1 - \pi_i)] < 1$. This property is called subcertainty and can help explain the Allais paradox. Expressing the Allais paradox in terms of prospect theory, we get:

$$v(100) > v(500) \cdot w(0.1) + v(100) \cdot w(0.89)$$

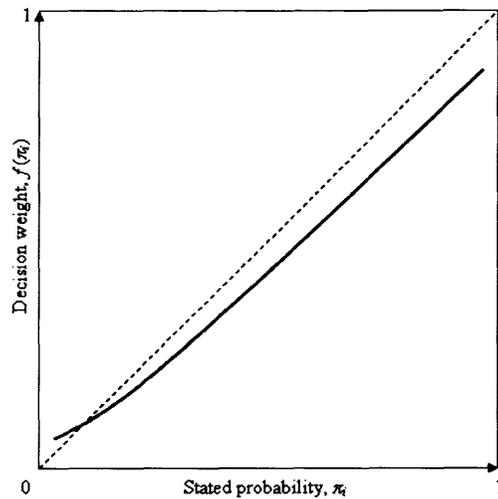
and
$$v(100) \cdot w(0.11) < v(500) \cdot w(0.1)$$

Therefore:

$$v(100) \cdot [1 - w(0.89)] > v(500) \cdot w(0.1)$$

and
$$v(500) \cdot w(0.1) > v(100) \cdot w(0.11)$$

Figure 2.6. Kahneman and Tversky's weighting function



Therefore:

$$v(100) \cdot [1 - w(0.89)] > v(100) \cdot w(0.11)$$

Therefore, it must be that

$$1 > w(0.89) + w(0.11)$$

Figure 2.6 illustrates a weighting function that fulfils all of the properties discussed.

2.2.3.3. Empirical assessment. In their updated version of prospect theory, Tversky and Kahneman (1992 [94]) provide parameters that describe the representative decision maker. Following much empirical testing, they propose the

following value function:

$$\begin{aligned} v(x_i - r) &= (x_i - r)^{0.88} && \text{for } x_i \geq r \\ &= -2.25(x_i - r)^{0.88} && \text{for } x_i < r \end{aligned}$$

where 2.25 is the coefficient of loss aversion, and 0.88 is the coefficient of risk aversion.

2.2.4. Criticisms of the Heuristics Paradigm

The heuristics paradigm began to receive significant criticism at the beginning of the 1980s. Some criticisms focused on the vagueness of the heuristics, their contradictory nature and the failure to specify when certain rules will be used in preference to others. Others questioned the scientific method of the empirical investigation, claiming that Kahneman and Tversky's methodology was based on devising often complex experiments which would lead to poor performance amongst subjects. Other criticisms are based on the normative model of rationality against which Kahneman and Tversky compare their empirical results. They claim that the heuristics literature prescribes a limited normative model which fails to reflect the complexity of the world as it is.

I describe below a number of the most significant attacks on the heuristics paradigm.

2.2.4.1. Birnbaum. One of the first significant criticisms of the heuristics school of thought came from Michael Birnbaum (1983). He explained how the famous cab problem could be accounted for in terms of optimal behaviour.

Tversky and Kahneman (1982 [92]) use the cab problem to illustrate the neglect of base rates in human inference. The problem is stated as follows:

A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data:

- (a) 85% of the cabs in the city are Green and 15% are Blue.
- (b) a witness identified the cab as Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colours 80% of the time and failed 20% of the time.

What is the probability that the cab involved in the accident was Blue rather than Green?

In empirical tests, where subjects are given the problem above and asked for an answer, the median and modal response is 0.80.

Tversky and Kahneman claim that the correct answer to the above question can be straightforwardly ascertained using Bayes' theorem:

$$\begin{aligned}
 \textit{Posterior odds} &= \textit{Prior odds} \times \textit{Likelihood ratio} \\
 \frac{\Pr(B \mid \text{"B"})}{\Pr(G \mid \text{"B"})} &= \frac{\Pr(B) \Pr(\text{"B"} \mid B)}{\Pr(G) \Pr(\text{"B"} \mid G)} \\
 &= \frac{0.15}{0.85} \cdot \frac{0.8}{0.2} \\
 &= \frac{12}{17}
 \end{aligned}$$

Therefore:

$$\Pr(B \mid \text{"B"}) = \frac{12}{12 + 17} = 0.41$$

where B and G denote the hypotheses that the cab involved in the accident was Blue or Green, and "B" is the witness's report that the cab was Blue. According to this result, the cab involved is more likely to be Green than Blue, despite the witness's report, because the base rate is more extreme than the witness is credible. Tversky and Kahneman claim that the divergence between their normative solution of 0.41 and the empirical responses is explained by subjects ignoring the base rates, and that this in turn can be explained by the representativeness heuristic.

On a purely intuitive level, it may seem strange that the report of a witness whom we are told is 80% reliable should carry so little weight. This is because the 'normative' solution here is based on the unrealistic assumption that the likelihood ratio is independent of the base rate. Tversky and Kahneman's calculations

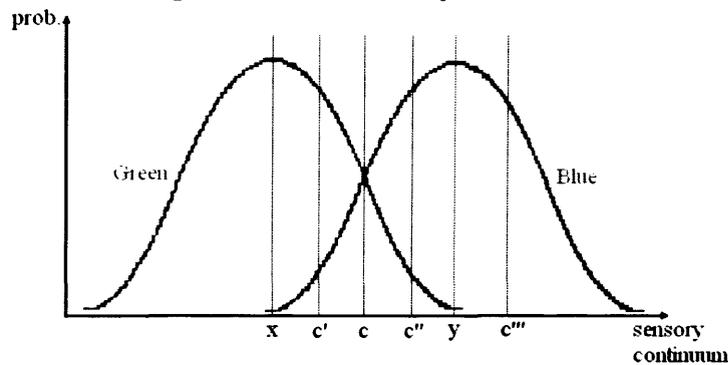
implicitly assume that $\Pr(B | "B") = \Pr(TS)$ and $\Pr("B" | G) = (1 - \Pr(TS))$, where $\Pr(TS)$ is the probability of success in the court test. In other words, they assume that the witness is equally likely to be wrong when reporting that he has seen a Blue cab as when he reports seeing a Green cab. Signal detection theory suggests that this is not the case. It suggests that the success of the witness in reporting each of the colours is different and, furthermore, the success rate in each case will vary as the base rates of the colours vary.

Figure 2.7 illustrates the theory of signal detection. Each colour of cab produces a normal distribution on a sensory continuum. For example, a Green cab is most likely to give sensation x but can, with tiny probability, give rise to sensation y . The witness will receive a sensation from the cab and, based on this, must then judge whether the cab is Blue or Green. As an example, the witness may choose sensation c as his judgement criterion. If the sensation lies to the left of c he will report that the cab is Green and if it is to the right of c he will report Blue. Tversky and Kahneman's normative analysis implicitly assumes that this judgment criterion is invariant, at c .

The key to Birnbaum's argument is that the judgement criterion that the witness uses will be dependent upon the base rate. There are three theories about how the judgment criterion is determined:

- (1) Range-frequency theory states that judges have a tendency to use response categories with equal frequency. Therefore, as Blue becomes less common in the base rate, the judgement criterion moves in favour of reporting Blue.

Figure 2.7. The sensory continuum



In other words, the judgement criterion moves towards the left, from c to c' . This makes the report of Blue more likely than under Tversky and Kahneman's analysis and so suggests a normative solution to the cab problem of less than 0.41.

- (2) Probability matching theory suggests that if the witness knows the base rate, it seems plausible that he will try to match response probabilities to stimulus probabilities, i.e. the witness will try to set $\Pr(\text{"B"}) = \Pr(B)$ and $\Pr(\text{"G"}) = \Pr(G)$. In this case the judgement criterion will shift right to c'' when Blue has a low base rate, so as to equate the probability of reporting Blue to 15%, the actual proportion of Blue in the population. This implies a normative solution of 0.58.
- (3) Clearly, the above solutions are both conjectures about human behaviour, but neither represents an optimal response to the problem that the witness faces. The optimal observer will adjust his judgement criterion so as

to maximise the probability of a correct identification. If $\Pr(B)$ is low, he will be more concerned about reporting Green correctly than getting Blue wrong. This implies a judgement criterion of c'''' and a normative solution to the cab problem of 0.82.

Intuitively, the lower the proportion of Blue cabs in the population, the more averse the witness will become to reporting Blue. He is only likely to report Blue when he is very confident that the cab is indeed Blue (when the probability of the sensation he experiences being caused by a Green cab is extremely remote). The subjects responding to Tversky and Kahneman's empirical studies will have taken this effect into account, and so, when the witness reports seeing a Blue cab, they recognise that he would only have done so if he were very confident that the cab was not Green, and it is this that allows them to have such confidence in the witness's report.

Birnbaum's criticism centres on the normative model used by Tversky and Kahneman. Their use of Bayes' theorem in the equations above assumes that the base rate only affects the prior odds of the possible outcome. The account of signal detection theory given above, however, makes clear that this is not so. The likelihood ratio is also dependant upon the base rate, because this affects the optimal response of the witness. Taking this into account, we find that the empirical modal response of 0.80 is very close to the optimal normative solution of 0.82.

2.2.4.2. Gigerenzer. Gerd Gigerenzer (1991 [46]), again, bases his criticism of Tversky and Kahneman's theory on their naïve interpretation of probability theory. The main thrust of his argument centres on the fact that many of Tversky and Kahneman's experiments require subjects to provide probabilistic interpretations of single events. According to the frequentist interpretation of probability this is not valid.

For example, Kahneman and Tversky often use experiments such as the following:

Which city has more inhabitants? (a) Hyderabad or (b) Islamabad

How confident are you that your answer is correct?

Typically, such experiments find that only around 80% of people who report being 100% confident in their answer are actually correct; only around 75% who report being 90% confident are correct; and so on. Kahneman and Tversky take this to be evidence of an overconfidence bias. But a discrepancy between confidence in a single event and relative frequencies in the long run is not a violation of probability theory from many points of view, because we are comparing two inherently different things. The frequentist school rejects the idea that 'probability' can even refer to a single event, whilst the subjectivist school only requires the internal consistency of subjective probabilities. Gigerenzer reports on an experiment in which he rephrased the above experiment in a frequentist manner. His subjects were asked fifty questions similar in type to the one above, and were then asked

how many they thought they had answered correctly. In this way, the reported frequency of correct answers could be compared to the actual frequency of correct answers. No evidence of overconfidence was found.

Gigerenzer similarly criticises the analysis of the conjunction fallacy. Tversky and Kahneman (1982 [93]) illustrate the conjunction fallacy using the ‘Linda’ experiment. In this experiment, the subjects were presented with the following brief personality sketch:

Linda is a 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

and asked to rank a number of statements in order of probability. The statements included, amongst others, the following:

- Linda is active in the feminist movement (F)
- Linda is a bank teller (T)
- Linda is a bank teller and is active in the feminist movement(F&T)

Tversky and Kahneman report that in a number of trials, with subjects of varying statistical ability, F&T is consistently reported as being “more probable” than T. This, they argue, violates the conjunction rule and is explained by the representativeness heuristic. Linda resembles more closely a “feminist bank teller” than a “bank teller” and so subjects judge the former to be more probable. Gigerenzer

argues that this experiment is just another example of subjects being asked for the probability of a single event, and so is meaningless from a frequentist view. We clearly still have a problem from a subjectivist viewpoint, in that the results are a clear breach of Bayes' theorem. However, Fiedler's (1988 [40]) rephrasing of the problem in a frequentist manner causes the conjunction fallacy to disappear and this result cannot be accounted for by the representativeness heuristic.

Finally, Gigerenzer criticises the analysis of base rate neglect. Again, Kahneman and Tversky's empirical studies almost always refer to single events. Even on a subjectivist reading, however, it is still unclear in most cases what the appropriate base rate is to use. Furthermore, as Birnbaum has made clear, a base rate only matters if subjects are drawn randomly from the population to which the base rate refers. Gigerenzer, Hell and Blank (1988 [47]) report a version of Kahneman and Tversky's (1973 [55]) engineer problem, an experiment in which they get subjects to draw descriptions from an urn in order to make clear the randomness. They find that the base rate neglect disappears. When Cosmides and Tooby (1996 [28]) rephrase the clinical trial problem (Tversky and Kahneman 1982 [92]) in a frequentist way and include extra information to explain the relevance of the base rate, they also find that the bias disappears.

2.2.4.3. Further criticisms. During the 1990s and 2000s a large amount of literature has been published which can be considered as a continuation and deepening

of the line of enquiry opened by Birnbaum. It focuses on the crucial role that the environment plays in shaping human behaviour. Important contributions in this area include Klayman and Ha's (1987 [61]) explanation of 'confirmation bias' and Oaksford and Chater's (1994 [72]; 1996 [73]) explanation of the 'selection task'.

Wason (1960 [96]) claims that the best way to test a hypothesis is to try to disprove it through negative testing. However, he finds that in empirical experiments his subjects almost always try to test a hypothesis through positive testing, testing which seeks to confirm the hypothesis. He refers to this anomaly as 'confirmation bias'. Klayman and Ha, however, show that the optimal strategy for testing hypotheses depends on the structure of the problem, and is not always to try and disconfirm the hypothesis. They argue that when hypotheses concern minority phenomena then positive testing is most effective. Furthermore, this is the most common type of hypothesis testing that occurs in the real world. Hence, whilst Wason can provide an example of a laboratory test in which subjects' use of positive testing is maladaptive, this behaviour is none the less highly adapted and optimal in most real life situations.

Wason (1966 [97]; 1968 [98]) provides the 'selection task' as an illustration of confirmation bias. The task involves four cards, each with P or $\neg P$ on one side and Q or $\neg Q$ on the other. One card shows P, another $\neg P$, another Q and the other $\neg Q$. The subjects are required to test the rule that "if P, then Q" by turning over as few cards as necessary. Typically, fewer than 10% choose the cards that standard logic dictates might falsify the hypothesis, P and $\neg Q$. The majority

of subjects choose to turn over P and Q. Oaksford and Chater argue that this response is not irrational. They show that Bayesian induction prescribes that P and Q are the most informative cards with respect to testing the hypothesis if P and Q are relatively rare compared to $\neg P$ and $\neg Q$. Again, performance in this task can be regarded as highly adaptive in a world in which we tend to test hypotheses about rare events. In this contrived context, the behaviour of subjects may not be optimal but it is likely to be optimal in the real world.

As with the cab problem, much of the behaviour observed in the laboratory environment can appear maladaptive and irrational. However, when we take into consideration the complex environment in which judgments must always be made, that same behaviour can appear compatible with the dictates of optimality.

2.2.4.4. Summary. Kahneman and Tversky's ([54],[55],[91],[92],[93]) methodology in their heuristics literature has been to list a range of biases that they have observed in their empirical work and then to explain these biases in terms of cognitive heuristics. Clearly, in order to differentiate between 'correct' and 'erroneous' judgments, it is necessary to have a normative model against which to compare actual judgments. In nearly all cases, the criticisms of the heuristics literature revolve around the normative model against which actual decisions have been measured. As Gigerenzer (1991 [46]) argues, they "have relied on a very narrow normative view and have ignored conceptual distinctions fundamental to

probability theory". Essentially, the heuristics paradigm stands accused of simplifying the theory of probability to the point of nonsense, and then concluding that actual decision making is biased because it fails to comply with this simplified probabilistic model.

2.3. Rationality and Irrationality in Economic Modelling

So far I have investigated the details of the two major competing theories of how agents make decisions when faced with uncertainty: subjective expected utility theory and the heuristics and biases paradigm. I have also considered some of the many critiques of these two theories.

As economists, our main concern in deciding between the competing theories of rational and heuristic agents is likely to be in the effect of these theories in the modelling of final market outcomes. If individual agents are sub-rational in ways that other agents in the market can correct for, then the correct specification of the decision making process may be relatively unimportant in modelling the macroeconomy³. Similarly, if the ‘mistakes’ that agents make are random and uncorrelated, then at the macro level they may net out. In this case rationality would, again, be a reasonable assumption on which to base macroeconomic models. Due to these issues, it makes sense to extend our models of rational and heuristic man to the level of the market so that we can draw comparisons between the competing theories at that crucial level.

The natural environment in which to challenge the two theories is in the relatively transparent setting of the financial markets. Here the role of preferences is likely to be less obfuscating than in other markets because of the relatively homogenous nature of financial products (a mix of risk and financial value). Because

³Of course, in this case sub-rationality may still have significant consequences in terms of the distribution of wealth, and so accurate modelling of decision making would be important in addressing distributional concerns.

of this, it is unsurprising that the bulk of work based on heuristical assumptions has been in the field of finance. Behavioural finance has provided a popular area for research over recent years, and I will provide an overview of the work in this field below. Before that, I will consider the theory against which the behavioural finance literature sets itself - the efficient markets hypothesis.

2.3.1. The Efficient Markets Hypothesis

The efficient market hypothesis (EMH) states that actual market prices of assets are equal to their fundamental values. This fundamental value is the discounted sum of expected future cashflows, where the expectation is taken over the correct distribution of possible outcomes conditional on currently available information. In an efficient market, 'prices are right' in that no investment strategy can earn excess returns. Higher returns are only available if higher risk is accepted.

The EMH is the natural conclusion of the theory of rational agents. No rational agent would choose to sell an asset for less than its fundamental value and no rational agent would choose to buy an asset for more than its fundamental value. Hence, the only price at which exchange can take place in a market populated by rational agents is that which is equal to the asset's fundamental value.

Given that a rational agent's expectation of future cashflows will be contingent upon all information available about past, present and future events, current prices must also take account of all the available information. Therefore, changes in price

expectation must simply reflect changes in the information set available from one period to the next. If everyone expects that prices in the next period will be £100, then our pricing decisions this period will reflect that expectation. Any changes in the price between now and then must either reflect the necessary rate of return on the asset (the difference between the price now and £100) or new information coming to light about the fundamental value of the asset (the difference between the price next period and £100). As Bachelier (1900/1964 [10]) puts it in the original discussion of the hypothesis, “past, present, and even discounted future events are reflected in market price, but often show no apparent relation to price changes”.

This implication of the EMH can be formally stated as follows:

$$X_t - E[X_t | \Phi_{t-\tau}] = \varepsilon_t \quad ; \quad E[\varepsilon_t \cdot \Phi_{t-\tau}] = 0$$

Prices, X , will vary from their expected values by a random variable ε , uncorrelated with the information set Φ available at the time that the expectation is formed. In other words, any differences between asset price movements and the expected movements implied by the market will be random and unforecastable.

Roberts (1967 [78]) posits three different forms of the efficient market hypothesis, each defined in respect to the information set that is prescribed. The weak version defines the information set Φ as all past asset prices. The semi-strong version defines Φ as all past and present public information, so that it includes such

things as announcements of annual earnings, stock splits, and so on. The strong form of the hypothesis defines Φ as all past and present information, public and private.

At first sight, it seems that the claim of unforecastability should be straightforward to test, but this is not the case. Whilst the price level at time t , X_t , and the information set at time $t - \tau$, $\Phi_{t-\tau}$, are observable; the expectation formed at time $t - \tau$ of the price level at time t , $E[X_t | \Phi_{t-\tau}]$, is not observable. In order to test the relationships described by the above equations we must, therefore, determine the price expectation implied by the market. This involves positing a theory of expected returns. The most popular theory of expected returns is the capital asset pricing model (CAPM).

2.3.1.1. The capital asset pricing model. A rational risk-averse investor will always require a risk premium on any risky asset. By combining assets in a portfolio he will be able to diversify away any unsystematic risk associated with the asset, but there will remain a certain amount of systematic risk which cannot be diversified away. This systematic risk is the extent to which the return on the asset moves with the market, the covariance between the return on the individual asset and the return on the market as a whole. The greater the systematic risk, the greater the return that the investor will require from the asset. In this way, an investor's asset portfolio can be seen as a trade-off between the expected return

and the variance of the return on the portfolio. Each investor will hold a mean variance efficient portfolio, a portfolio which gives maximum expected returns for a given variance.

Given that rational investors, with the same information set, will have the same beliefs about expected returns and covariances of individual assets, the market portfolio will also be mean variance efficient. It can be shown that the expected returns on individual assets will then be linearly related to the expected return on the market portfolio:

$$(2.1) \quad E[r_{jt} - r_{ft}] = \beta_j \cdot E[r_{mt} - r_{ft}]$$

Where r_{jt} is the return on asset j in period t , r_{mt} is the return on the market in period t , and r_{ft} is the return on riskless assets such as treasury bills. β_j is a measurement of the systematic risk of asset j . It measures the extent to which variations in the returns on asset j are related to the fluctuations in the return on the market as a whole, and is given by:

$$\beta_j = \frac{Cov[r_{jt}, r_{mt}]}{Var[r_{mt}]}$$

Investors are compensated for bearing this systematic risk, $E[r_{mt} - r_{ft}] > 0$, because they cannot eliminate it through diversification of their portfolio.

2.3.1.2. Consumption CAPM. Lucas' (1978 [64]) representative agent theory of asset pricing yields the following first-order condition for the optimal consumption and portfolio decision:

$$u'(c_t) = E [\rho \cdot (1 + r_{jt}) \cdot u'(c_{t+1})]$$

where c_t is consumption in time period t and ρ is the time preference of consumption. This result yields the same expected return to holding asset j as in equation (2.1) above, but where β_j is as follows:

$$(2.2) \quad \beta_j = \frac{Cov[r_{jt}, \gamma_{t+1}]}{Cov[r_{mt}, \gamma_{t+1}]} \quad ; \quad \gamma_{t+1} = \frac{\rho u'(c_{t+1})}{u'(c_t)}$$

Equation (2.2) states that an asset's expected return must be higher the more negative the covariance between the asset and the ratio of marginal utilities. The intuitive reason for this is that when consumption is low, marginal utility is high and γ_{t+1} is large. An asset with a more negative covariance with γ_{t+1} , therefore, has lower returns when consumption is low, which is precisely when wealth is more valued. An investor would therefore require a higher premium to hold such an asset.

2.3.1.3. Empirical Evidence of Efficiency. By combining the EMH and CAPM, we can draw several conclusions about the outcomes of asset markets populated by fully rational agents, in the SEU sense.

Firstly, ‘prices are right’ implies that no investment strategy can earn excess returns. Chartists, who predict future price movements by seeking to interpret past patterns, will fail to do any better than a market tracking portfolio. Similarly, fundamental analysts who study corporate reports to try and gain insights into the real worth of shares, will fail to earn an excess return because shares will already reflect fundamental value. In short, beyond diversifying away unsystematic risk, fund management is not profitable.

Secondly, consumption CAPM implies that stock returns should be correlated with consumption growth. Therefore, returns may be predictable to some extent, but only to the extent that consumption is predictable.

As Fama (1970 [36]) notes in his wide ranging review of the literature on efficiency, much empirical work in the field of efficient markets preceded any development of the theory. As early as 1900, Bachelier concluded that commodity prices fluctuate randomly, and further studies by Working (1934 [99]) and Cowles and Jones (1937 [29]) showed that US stock prices and other economic series follow the same pattern. Kendall (1953 [59]) examined 22 UK stock and commodity price series and concluded that “the data behave almost like wandering series”.

All of this predates the prescription of a formal theory of efficient markets, which has its origins in Samuelson (1965 [81]). He derives a general model of fair-game futures pricing which demonstrates that there is no way of making an expected profit by extrapolating past changes in the future’s price.

Fama finds that results are strongly supportive of the weak form of the hypothesis, and that the semi-strong version is also supported. The strong form does not appear to be supported, and Fama suggests that such an extreme model should be regarded as a benchmark rather than a descriptive model. He concludes by remarking that evidence in contradiction to the efficient markets model for price changes covering periods longer than a single day is hard to find.

More recently, however, Fama's conclusions have been challenged widely. Lo and MacKinlay (1988 [62]) find serial correlation in stock price changes using weekly US data for 1962-1985, though they cannot reject the random walk model on the basis of a four week interval. Other studies, including Fama and French (1988 [37]), have shown evidence of negative serial correlation (mean reversion) over longer time periods. These phenomena have been labelled as the 'predictability puzzle' by Barberis and Thaler (2002 [14]).

Mehra and Prescott (1985) use a variation of Lucas' (1978) pure exchange model to examine US financial data used the period 1889-1978. They find that the data strongly violate restrictions that general equilibrium models impose upon the average returns of equity and Treasury bills. The explanation of the excess return on equities over risk free investments⁴ implies an unfeasibly large coefficient of risk aversion. This result is known as the equity premium puzzle.

⁴Since 1926 the annual real return on US stocks has been about seven percent, versus a real return on treasury bills of less than one percent (Benartzi and Thaler 1995 [19])

Other anomalies in the data that need to be explained include the high volatility of equity returns (for example, Campbell and Cochrane 1999 [25]); the negative long run performance of new issues (Ritter 1991 [77]); the outperformance of larger stocks by smaller stocks (Banz 1981 [11]); and various seasonal effects, including month of the year, day of the week and hour of the day effects (French 1980 [44]; Harris 1986 [52]; Rozeff and Kinney 1976 [80]).

2.3.1.4. Development of the EMH. As discussed earlier, the EMH has no practical relevance unless it is coupled with a theory of expected returns. We must posit some underlying asset pricing model, such as the CAPM, if we wish to test the hypothesis of market efficiency. However, any such test will consequently become a joint test of the EMH and the relevant asset pricing model. The puzzles discussed in the previous section, therefore, do not necessarily require us to revise or discard the notion of market efficiency. Rather, it may be the case that we require a more sophisticated model of asset pricing.

Since the late 1980s researchers have made various attempts to develop more sophisticated asset pricing models than the CAPM and consumption CAPM in order to try and capture the effects seen in the data. Some of the literature has been based on using ever more sophisticated specifications of utility (for example, Abel 1990 [1]). Others have modelled variable risk aversion (Campbell and Cochrane 1999 [25]). A major focus for research is the ‘Peso problem’, which suggests that

the variability in a sample may be unrepresentative of the true long run population variability. This is likely to occur when there is a small probability of a catastrophic event. Agents will take this into account in their expectation formation, but the event is very unlikely to have an effect on observed outcomes. For example, Rietz (1988 [76]) specifies a model in which there is a small probability (0.4% or 1.4%) of a massive fall in output (50% or 25%). In such a situation, risk-averse investors demand a large risk premium on equity to compensate for the extreme losses that will occur in the case of an unlikely but severe market crash. Rietz concludes that this can explain the equity premium puzzle.

2.3.2. Behavioural Finance

As we have already remarked, the traditional finance paradigm, in which agents are assumed to be rational, has struggled to explain a number of phenomena that have been observed in real world financial markets. These phenomena include the equity premium puzzle, the predictability puzzle, and the volatility puzzle. Behavioural finance attempts to explain these occurrences by using models in which some agents are not fully rational.

The rationality of agents in traditional models is two dimensional. In the first place, agents' beliefs are assumed to be correct. The subjective distribution that they use to forecast future realisations of unknown variables is indeed the distribution that those realisations are drawn from. This is the rational expectations

hypothesis. Secondly, the agents are assumed to make optimal decisions given their beliefs. In other words, they maximise subjective expected utility.

Behavioural finance seeks to analyse what happens when we relax either of these two assumptions. In some behavioural finance models agents do not hold rational beliefs, whilst in others they hold correct beliefs but make choices which are suboptimal in terms of subjective expected utility theory.

It is not enough, however, merely to explain why some people may behave irrationally. Behavioural finance must also explain how any rational agents in the economy are prevented from arbitraging prices to their 'rational' level. Explaining how, in an economy in which irrational and rational agents interact, irrationality can have a long term impact on prices is one of the two pillars of the behavioural finance literature (the other is the explanation of irrationality provided by the psychology literature). We will therefore begin our discussion of behavioural finance with a justification of how arbitrage may be limited in real world markets.

2.3.2.1. Limits to Arbitrage. Behavioural finance argues that asset prices are sometimes best interpreted as deviations from fundamental value, brought about by the presence of irrational traders in the economy. The straightforward response to this from efficient market theorists is that rational traders will rapidly undo any mistakes made by the irrational traders.

John Kay (2004 [58]) draws an analogy between the efficient market and the queues at a supermarket. Rational agents at the supermarket checkout will choose to stand in the queue with the shortest expected waiting time. In a world solely inhabited by rational agents, this will cause the equilibration of waiting times at each of the checkouts – without any explicit coordination, the queues in the supermarket will be of roughly the same length. However, even in a world where many of the agents do not optimise by choosing the shortest queue, but rather choose any queue randomly, the existence of a relatively small number of rational agents will ensure the same result as before. These rational agents will all move towards the shortest queues, and this will cause an equilibration of queue length just as before. The rational agents will gain more from their optimising behaviour than they would have done in a fully rational world, but the aggregate result is as before. Similarly, in a stock market, irrational traders (noise traders) who become overly pessimistic about a company may undervalue its shares at £15 even though its fundamental value is £20 a share. However, even if there is just a single rational agent (arbitrageur) in the market, he will sense an opportunity to buy shares in the company at their bargain price and will continue to do so until the price is driven up to its fundamental value.

Arbitrage is often assumed to be a riskless activity. If there is a mis-pricing, it is assumed that a riskless profit is there for the taking because the price can be forced back to its fundamental value. Behavioural finance argues that a mis-pricing does not always provide an attractive investment opportunity, because corrective

strategies can be so risky as to make them unattractive to the rational investor. In this way, arbitrage may be hindered and the price of an asset may not reflect its fundamental value. If there is a 'free lunch' to be had, any rational investor will take it, but behavioural finance theory aims to explain why a mis-pricing does not always provide the opportunity for a 'free lunch'.

Barberis and Thaler (2002 [19]) list four types of risk faced by a potential arbitrageur:

- (1) Fundamental risk: an arbitrageur faces the risk that, having bought an underpriced stock, the fundamental value of the stock falls in the light of subsequent bad news. This risk can be hedged by shorting a substitute security. However, substitutes are rarely perfect and so fundamental risk can never be completely removed.
- (2) Noise trader risk: the mis-pricing that is being exploited by the arbitrageur may worsen in the short run. The irrationally pessimistic investors who caused the mis-pricing in the first instance, may become even more pessimistic and cause the asset price to move even further away from its fundamental value. Noise trader risk will be irrelevant to arbitrageurs with infinite horizons, but arbitrageurs do not have infinite horizons. In fact, they usually have very short horizons because of liquidity constraints. They are usually fund managers, managing other people's money. Agency problems will require that returns are positive in the short as well as the

long run, because investors are likely to judge the performance of the fund manager on the basis of short term returns.

- (3) Implementation costs: commissions and bid-ask spreads will obviously limit arbitrage. However, legal constraints may also come into play. For example, many fund managers are simply not allowed to engage in shorting, and this may prevent them from hedging against fundamental risk.
- (4) Model risk: the arbitrageur may have assessed the fundamental value of the asset wrongly and a perceived mis-pricing simply does not exist.

The key to arbitrage is that rational traders can make a sure profit when an asset is mis-priced, because in the long run asset prices must converge towards their fundamental value. However, if we allow for the fact that a mis-pricing can occur in the market, we must also allow that in the short run this mis-pricing could well increase. If rational traders face a liquidity restraint then they cannot be sure that when they need to liquidate their holdings of an underpriced asset, that the underpricing will have been corrected. De Long et al (1990 [32]) present a model which demonstrates that “as long as arbitrageurs have short horizons and so must worry about liquidating their investment in a mis-priced asset, their aggressiveness will be limited even in the absence of fundamental risk”. The addition of fundamental risk and other practical issues such as implementation costs merely exacerbate the riskiness of arbitrage, and make the case for complete arbitrage less persuasive.

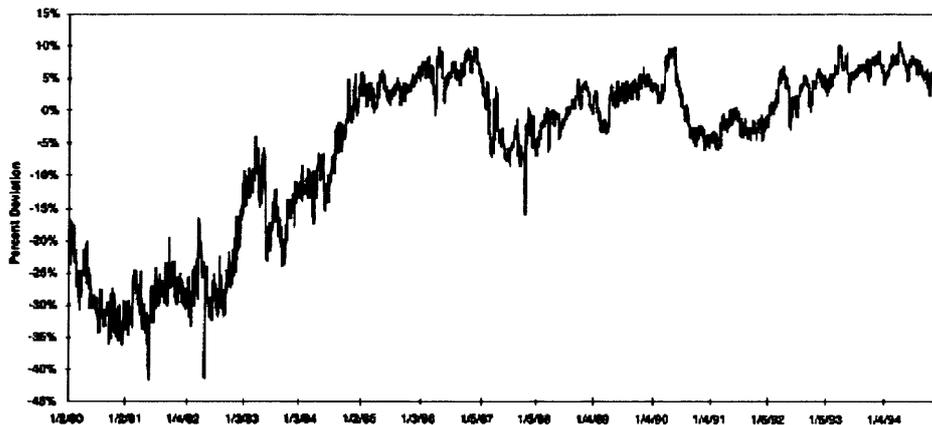
It is difficult to establish the existence of a mis-pricing in the real world because to do so normally involves having to model the fundamental value of an asset,

and any such model process will obviously be open to criticism. There are some instances, however, in which a mis-pricing can be observed without having to refer to fundamental values. Barberis and Thaler (2002 [14]) provide several examples of where arbitrage can be seen to be deficient, including ADRs, index inclusions and internet carve-outs. We will illustrate the possibility of mis-pricing with the example of twin shares.

The Royal Dutch / Shell group has shares listed on several stock markets. Royal Dutch shares have a claim to 60% of the total cashflow of the group, and are traded primarily in the US and in the Netherlands. Shell shares have a claim to 40% of the total cashflow of the group, and are traded primarily in the UK. If share prices equal fundamental value, then the value of Royal Dutch equity should always be $1\frac{1}{2}$ times the value of Shell equity. Figure 2.8, taken from Froot and Dabora (1999 [45]), illustrates how the ratio of Royal Dutch to Shell equity value has deviated wildly from the ratio of $1\frac{1}{2}$.

By buying the relatively underpriced stock and shorting on the relatively overpriced stock, an arbitrageur could exploit the mis-pricing whilst being perfectly hedged against fundamental risk (in fact we have not even had to refer to fundamental value in determining that a mis-pricing exists). Clearly, the only thing holding back potential arbitrageurs is noise trader risk. Although they know that the fundamentals must in the long run drive the equity ratio back to $1\frac{1}{2}$, the evidence from figure 2.8 suggests that the mis-pricing can be persistent.

Figure 2.8. Deviations in Royal Dutch / Shell equity values



2.3.2.2. Modelling an Irrational Market. As stated above, behavioural finance can be viewed as the relaxation of the assumptions of rationality inherent in traditional finance theory. Clearly, the removal of the straightjacket of rationality has the potential of allowing us to provide an explanation of phenomena which cannot be explained by traditional models. Much has been written on an intuitive or anecdotal level about how certain beliefs might explain some of the puzzles of financial data. For example, Barberis, Schleifer and Vishny (1998 [13]) claim that conservatism suggests that people put too little weight on the latest piece of earnings news relative to their prior beliefs. Their model, based on a complex belief structure, can explain such features of the data as post-earnings announcement drift. Daniel, Hirshleifer and Subrahmanyam (1998 [30]), on the other hand, stress biases in the interpretation of private information. They claim that overconfidence in private information generates high volatility and long term mean reversal.

It can appear that there is a tendency in the literature to develop models that will explain observed phenomena by construction, rather than to develop models based on the theory and then test that against the data. The two models mentioned above make very different assumptions about investors' decision making and are able to explain some, but certainly not all of the anomalies in the data.

We consider below three further models which are constructed based on various different assumptions about decision making. The first is based directly on prospect theory. The second represents an attempt to incorporate some of the features of prospect theory into a general equilibrium framework. The third is based on the idea that agents employ simple heuristical rules.

Benartzi and Thaler (1995 [19]) attempt to provide an explanation of the equity premium puzzle by developing a model in which agents are loss averse, as prescribed by prospect theory (Kahneman and Tversky 1979 [56]). Specifically, they prescribe a value function of the following form:

$$(2.3) \quad \begin{aligned} v(x-r) &= (x-r)^\alpha && \text{for } x \geq r \\ &= -\lambda(x-r)^\beta && \text{for } x < r \end{aligned}$$

where $\lambda > 1$ is the coefficient of loss aversion. The parameters α and β are both specified as 0.88 and λ as 2.25, in accordance with the estimations of Tversky and Kahneman (1992 [94]). One of the consequences of using prospect theory is that

the frequency with which returns are evaluated must be specified. In prospect theory, value is determined in relation to changes in wealth relative to an initial asset position, r . The frequency with which this reference point is updated is a key determinant of the value function. In a model with loss aversion, the more often an investor evaluates his portfolio, and updates the value of r , the less attractive he will find a high risk investment. Benartzi and Thaler, therefore, ask the question, if investors have preferences as defined by equation (2.3) above, how often would they have to evaluate their portfolios in order to explain the equity premium puzzle?

In order to answer this question, Benartzi and Thaler draw samples from 1926-1990 monthly returns on stocks, bonds and treasury bills and use simulation techniques to compute the prospective value of holding stocks and bonds for various evaluation periods. A further complication is added by having to choose whether to use nominal or real returns. Clearly, losses will be more prevalent with each asset type when considering real returns. The results for nominal returns on stocks and bonds are shown in figure 2.9. The graph shows that the prospective value of stock and bond portfolios is equal when evaluated at thirteen month intervals. For real returns, the comparable result is between ten and eleven months. Benartzi and Thaler conclude that such a time frame is highly plausible, given that investors file tax returns annually and receive the most comprehensive information about the performance of their investments on an annual basis.

This result is, however, fairly meaningless in that the decision faced by investors is not between a portfolio of bonds and a portfolio of stocks. Rather, investors are

Figure 2.9. Prospective value of holding pure stock and pure bond portfolios

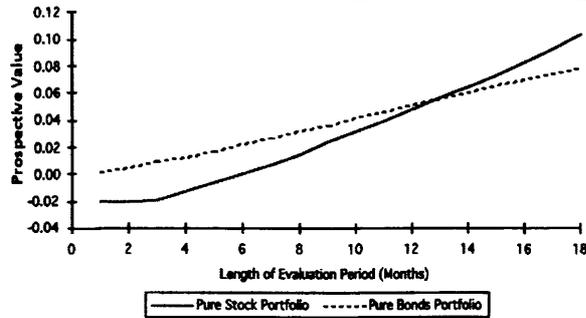
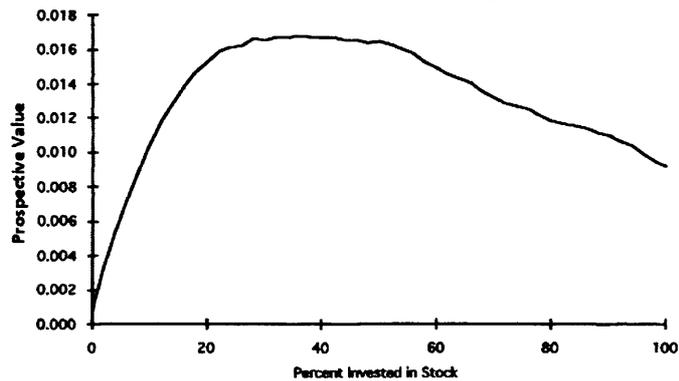


Figure 2.10. Prospective value of holding mixed portfolios



faced with the decision of what proportion of bonds and stocks to hold in their portfolio. Benartzi and Thaler perform a second simulation exercise which assumes the one year evaluation period which they claim to have established and then searches for the optimal mix of stocks and bonds. The results of the simulations using nominal returns are shown in figure 2.10. Portfolios with between 30 and 55 percent stocks yield the highest prospective value. This, the authors conclude, is roughly consistent with the observed behaviour of investors. The equity premium

puzzle disappears, because the premium can be explained in terms of loss aversion and annual re-evaluations of asset position by investors.

Barberis, Huang and Santos (2001 [12]) make the first attempt to try and build prospect theory into a general equilibrium model of the stock market. They do not use an explicit value function in the manner of Benartzi and Thaler (1995 [19]), but instead augment a standard SEU utility function with a term which reflects changes in financial wealth. The utility function is specified as follows:

$$E \left[\sum_{t=0}^{\infty} \left(\rho^t \frac{c_t^{1-\gamma}}{1-\gamma} + b_t \rho^{t+1} v(X_{t+1}, S_t, z_t) \right) \right]$$

The first term is a standard risk averse utility over discounted consumption, with discount factor ρ and risk aversion parameter $\gamma > 0$. The second term represents utility from fluctuations in the value of financial wealth. X_{t+1} is the agent's financial gain or loss between time t and $t + 1$ (measured annually), S_t is the value of his risky assets at time t , and z_t measures his gains or losses prior to time t as a fraction of S_t . In other words, the representative investor gains utility not only from consumption but also from changes in the value of his financial wealth.

The function $v(\cdot)$ is discontinuous at the origin of X_{t+1} , and is steeper for negative values of X_{t+1} . Essentially, $v(X_{t+1}, S_t, 1)$ is the same as the specification of $v(\cdot)$ in equation (2.3) above. This implies an investor who is loss averse, in line with the prescriptions of prospect theory. The inclusion of the term z_t as a

determinant of $v(\cdot)$ implies that, despite the annual re-evaluation of the reference point against which gains and losses are measured, the investor has a long memory and does not completely disregard events which occurred prior to the last re-evaluation of his financial position. His risk aversion increases with prior losses and decreases with prior gains; an investor who is burned by an initial loss becomes more sensitive to additional setbacks whilst an investor who has made substantial gains becomes less wary of risk. This concept has been labelled as the 'house money' effect in the psychology literature.

Essentially, it is this variation in risk aversion, depending on prior stock market performance, which allows the model to generate stock returns which have a high mean, are highly volatile, significantly predictable, and not highly correlated with the underlying low variance consumption growth process. At the same time, the model generates a low and stable riskless return. Intuitively, an unexpectedly high dividend raises prices, as predicted in any efficient market model. However, in this model, the price increase also makes the investor less risk averse, and this makes the risky asset relatively more attractive, driving prices still higher. A similar story holds for negative dividend news. In this way, returns are substantially more volatile than the underlying dividend growth.

In a series of papers, Frankel and Froot (1986 [42], 1990 [43]) develop the idea that exchange rate dynamics can only be explained effectively if we posit that investors have heterogeneous expectations.

They show how the existence of different forecasting techniques can lead to excess volatility in exchange rates. Furthermore, they present evidence that such heterogeneity in forecasting does exist in real world financial markets. They present evidence that the use of chartist forecasting techniques is widespread. However, at the same time, other agents in the foreign exchange markets make use of models based on fundamentals to forecast future exchange rates.

Their data also shows that, over time, investors alter the forecasting techniques that they use. During the period of study (1978-1988) 'technical' analysis became much more prevalent within their sample of forecasting firms, at the expense of fundamental analysis. Indeed, Frankel and Froot conclude that the shift over time in the weights that the market places on different forecasting techniques is essential in explaining the large fluctuations in exchange rates witnessed during the 1980s.

De Grauwe and Grimaldi (2006 [31]) develop these ideas of a tension between fundamental and chartist forecasts into a fully specified model of the foreign exchange markets. They claim that they are able to use their model to replicate many of the characteristics of exchange rate time series, including high levels of kurtosis.

2.4. Behavioural Considerations in Macroeconomics

So far, I have concentrated on the main area of behavioural influence in economics, that is finance. A common perception is that behavioural considerations have not penetrated much further than this into economic theory.

Certainly, the direct influence of the heuristics and biases paradigm has been overwhelmingly focused in finance. However, other areas of economics, and macroeconomics in particular, have long been informed by psychology.

Many people consider Keynes' General Theory as the origin of macroeconomics as a distinct field of study. This work drew heavily on psychological concepts such as cognitive bias, herding and the pursuit of social status.

Keynes' appeal to "animal spirits" as the basis of investment decisions is well documented. However, he makes an equally strong appeal to psychological phenomena as the basis for his consumption function:

These eight motives [to save] might be called the motives of Precaution, Foresight, Calculation, Improvement, Independence, Enterprise, Pride and Avarice; and we could also draw up a corresponding list of motives to consumption such as Enjoyment, Short-sightedness, Generosity, Miscalculation, Ostentation and Extravagance. (*General Theory* chapter 9.I [60])

Keynes' intuitive understanding of psychology is certainly primitive and rudimentary, but he makes use of it in explaining many important economic phenomena. Foremost amongst these, perhaps, is the resistance of workers in accepting nominal wage cuts.

2.4.1. Money Illusion

Irving Fisher wrote of money illusion as long ago as 1928 [41], and the issue of price and wage stickiness has remained central to the study of macroeconomics ever since. Without it, explanations of the non-neutrality of money are unconvincing⁵. Although the psychological basis for money illusion is rarely discussed, nominal rigidities provide the basis for the entire New Keynesian literature.

Shafir Diamond and Tversky (1997 [83]) provide a wealth of evidence to support the idea that people tend to think in terms of nominal prices rather than real (or relative) prices. They “interpret money illusion as a bias in the assessment of the real value of economic transactions, induced by a nominal evaluation”. In other words, people are predisposed to assess different options within the frame of nominal currency rather than in a real frame. This is due to the salience of the nominal frame. Essentially, it is easy to assess nominal prices, whereas assessing

⁵New Classical explanations of non-neutrality, based on incomplete information, can only explain non-neutrality in the very short run.

real prices takes more effort. Nominal prices are a heuristic for assessing fundamental value - as a rule of thumb nominal prices are a convenient way of choosing between alternatives. It is only when this rule of thumb becomes ineffective, for example in times of hyperinflation, that people will resort to the costly process of calculating real prices.

Akerlof and Yellen (1985 [5],[6]) show that the losses associated with using this heuristic is relatively small. In a model with efficiency wages and monopolistic competition, they show that the cost to firms of failing to readjust prices in the face of a money supply shock are second order. The aggregate impact on output, however, is first order. This result is complemented by Fehr and Tyran's (2001 [39]) finding that in a model with both rational agents and agents who suffer money illusion, there is an incentive for the rational agents to imitate the behaviour of those who are deluded. In this way, acting in accordance with the nominal heuristic may have very little cost to individuals but at the same time it has a significant macroeconomic effect.

As I've already alluded to, there are relatively few papers which consider the psychological basis of money illusion compared to the number that assume price stickiness as a given. The majority of New Keynesian models make use of well established constructs such as Taylor or Calvo contracts. Such constructs are often criticised as ad hoc, and defences of their use rarely go beyond citing their effectiveness in reproducing empirical phenomena. However, their true justification

is in the fact that they are a parsimonious way of introducing complex psychological phenomena into macroeconomic models.

2.4.2. The Labour Market

A serious problem for the rational expectations literature has been to explain the existence of involuntary unemployment. In a world of free and efficient markets workers should always be able to find work as long as they are willing to accept a low enough wage. According to New Classical models, employment falls during recessions because an unexpected decline in aggregate demand reduces nominal wages, which workers mistake for a fall in the real wage and they therefore withdraw from the labour force. Such a ‘voluntary unemployment’ interpretation of recession simply doesn’t square with the empirical facts. The New Classical story implies that voluntary resignations should be counter-cyclical, whereas in reality they fall dramatically during recessionary periods.

A series of papers appeared during the 1970s and 1980s⁶ which put forward various efficiency-wage explanations to the observation that excess supply of labour does not seem to lead to aggressive wage cutting. The key insight of these models is that, because positive wage differentials lead to increases in the marginal productivity of labour, profit maximising firms have an incentive to pay a real wage above the market clearing level.

⁶See Yellen (1984 [100]) for a survey.

There are a number of explanations as to why higher wages may be associated with higher productivity. Some of these explanations are compatible with rational optimisation. One explanation is that paying a wage above the market clearing level provides an incentive for workers not to shirk (e.g. Shapiro and Stiglitz 1984 [84]). This view is compatible with the Marxist conception of unemployment as a labour disciplining device. Another explanation is based on transactions costs (e.g. Stiglitz 1974 [87]). In this case, setting wages above the market clearing level reduces costly labour turnover. Asymmetric information can also provide justification for setting high wages (e.g. Stiglitz 1976 [88]). If workers are of heterogeneous ability, and reservation wages are positively correlated with that ability, then adverse selection explains why higher wages and higher productivity coincide.

The problem with all of these explanations for efficiency wages is that there is often an alternative, pareto-superior solution to the problem identified. In the case of shirking, fees for starting employment provide the same discipline as involuntary unemployment but allow the market to clear efficiently (Eaton and White 1982 [33]). Similarly, fees paid by workers to cover transactions costs would allow for efficiency. In the case of adverse selection, the usual solutions of self-selection and screening devices may lead us to conclude that it is unlikely to account for much involuntary unemployment. More significantly, all of the explanations considered so far explain why involuntary unemployment persists, but they do not explain

its counter-cyclical behaviour. Why any of the issues discussed should be less significant during boom periods is unclear.

For all of these reasons, psychological explanations of efficiency wages are the most convincing. Reciprocity, fairness and adherence to group norms all provide coherent explanations of the empirical facts. Akerlof's (1982 [3]) gift exchange model shows how firms may improve the work norms of their workers by paying them in excess of the minimum required to keep them in the job. Fehr and Gächter (2000 [38]) provide evidence of the importance of reciprocal behaviour in employment contracts. All of these theories are based on ideas of equity and fairness that can be traced back to Adams (1963 [2]). In a similar vein, insider-outsider models rely on the assumption that insiders are able to coordinate so as to sabotage the inclusion of new workers. This can be explained by group formation theory (e.g. Roy 1952 [79]).

2.4.3. Inter-temporal Discounting

A commonly held view is that people tend to display myopia when it comes to choosing between current and future consumption. As a result, almost all states in the developed world provide some sort of income support for elderly citizens, and have tax systems which are intended to incentivise saving. Of course, in the New classical world of rational, optimising agents such schemes would be unnecessary

as agents, by assumption, choose the level of savings that maximise their lifetime utility. Given that undersaving cannot be explained by conventional market failures, such as externalities, we must either accept that undersaving does not exist, or that conventional models fail to account for an empirical phenomenon.

The evidence suggests that there is a significant difference between the inter-temporal discount rates that people think they ought to use and those that they actually do use. We could treat this as any other difference between reported and revealed preferences. Akerlof (2002 [4]), however, argues that the difference comes from the fact that welfare is given by a particular function, but people actually maximise a different function altogether. The difference between the two functions is driven by the salience of present consumption and a lack of self-control. This view is supported by empirical evidence that people use high discount rates to choose between immediate consumption and future consumption, but that they use a lower discount rate when the same choice is presented over an equidistant period at some point in the future. Loewenstein and Prelec (1992 [63]) show that a hyperbolic discount function can account for such behaviour.

Thaler and Sunstein (2008 [90]) report a wealth of evidence that suggests that when people have the opportunity to pre-commit to certain future choices, they are able to avoid these problems of inter-temporal inconsistency. For example, they report that when pension schemes are set up so that people can commit to saving out of future wage increases, savings rates increase significantly.

2.4.4. Asset Markets

In the previous section, I discussed in some depth the evidence for less than rational behaviour in financial markets. It suffices to note here that asset markets are central to any economic system. Hence, any serious attempt at explaining the macroeconomy must incorporate the lessons learnt in financial research.

2.4.5. Summary

It is clear from the foregoing discussion that psychological phenomena have informed a number of different areas of macroeconomic research, in particular financial economics. However, much of the core of macroeconomics research remains unconcerned with behavioural issues. The New Classical school of macroeconomics has been consistent in its adherence to the rational paradigm. The New Keynesian school, on the other hand, relies upon nominal rigidities to explain macroeconomic dynamics. But, even though such rigidities must arise from some deviation from rational behaviour, the cause of this is rarely made explicit. There remain large areas of macroeconomics where either there has been no attempt to develop beyond the rational paradigm, or else the psychological underpinnings are vague.

In the next two chapters I attempt to address two important issues that have yet to be posed by the behavioural macroeconomic literature. The first is whether

simple heuristical models are able to capture the dynamics of asset prices in ways that rational models fail. The second is whether such models can help to inform monetary policy.

CHAPTER 3

A Behavioural Explanation for Asset Prices

As discussed in section 2.3.1.3, the use of rational expectations (RE) modelling in the field of asset pricing has received much empirical criticism. The efficient markets hypothesis (EMH) gives rise to a number of well known puzzles in the data. These include the equity premium puzzle¹, various predictability puzzles², the volatility puzzle³, and various seasonal effects⁴. What is clear is that standard RE-EMH models fail to replicate many of the characteristics of real world asset price time series. Standard models would suggest that returns to holding assets should be normally distributed, whereas in the real world the distribution displays much higher levels of volatility and kurtosis (fat tails).

There have been a plethora of responses in the finance literature to these puzzles. Some of these responses have been compatible with the axiomatisation of man as a rational agent, whilst others require some departure from the normative model.

¹Mehra and Prescott (1985 [69]) find that the explanation of the historical excess return on equities over risk free investments implies an unfeasibly large coefficient of risk aversion.

²Lo and MacKinlay (1988 [62]) find serial correlation in stock price changes using weekly US data for 1962-1985, whilst Fama and French (1988 [37]) show evidence of negative serial correlation (mean reversion) over longer time periods.

³For example, Campbell and Cochrane (1999 [25]).

⁴These include month of the year effects (French 1980 [44]), day of the week effects (Harris 1986 [52]) and hour of the day effects (Rozeff and Kinney 1976 [80]).

At the rational end of this spectrum is research based on the ‘Peso problem’. This line of research is based on the fact that the variability in a sample may be unrepresentative of the true long run population variability. This is likely to occur when there is a small probability of a catastrophic event. Agents will take this into account in their expectation formation, but the event is very unlikely to have had an effect on observed outcomes. For example, Rietz (1988 [76]) specifies a model in which there is a small probability (0.4% or 1.4%) of a massive fall in output (50% or 25%). In such a situation, risk-averse investors demand a large risk premium on equity to compensate for the extreme losses that will occur in the case of an unlikely but severe market crash. Rietz concludes that this can explain the equity premium puzzle.

Using a similar specification, Meenagh, Minford and Peel (2007 [68]) show that they can account for the FTSE all-share index simply as the discounted sum of the rational expectation of future profits.

The second type of response to the puzzles mentioned above has been to develop more sophisticated asset pricing models which involve specifying the utility function in non-standard ways. Abel (1990 [1]) finds that he can derive equity premia which are as large as the historically observed equity premium in the United States when he specifies a non-standard utility function. In his specification, utility depends upon habit and on the consumer’s level of consumption relative to the lagged cross-sectional average level of consumption. Campbell and Cochrane (1999 [25]) similarly show how habit persistence in the utility function can cause

risk aversion to vary over the cycle, and how this can explain some of the volatility in stock returns. Both these papers make use of standard optimisation techniques but are not truly consistent with the RE-EMH because they introduce concepts of habit persistence and relative consumption which are inconsistent with the rational paradigm. Habit persistence and concern for social status are both well established in the psychology literature, but are inconsistent with standard conceptions of rationality.

The third type of response to the puzzles discussed above is more firmly based upon psychological foundations. This is known as the behavioural finance literature, and it discards the assumption of rational expectations in favour of agents who use heuristics, or rules of thumb, as the basis for decision making.

These models have often been based on the heuristics and consequent biases identified by Tversky and Kahneman (1974 [91]). For example, Barberis, Schleifer and Vishny (1998 [13]) claim that conservatism suggests that people put too little weight on the latest piece of earnings news relative to their prior beliefs. Their model, based on a complex belief structure, can explain such features of the data as post-earnings announcement drift. Daniel, Hirshleifer and Subrahmanyam (1998 [30]), on the other hand, stress biases in the interpretation of private information. They claim that overconfidence in private information generates high volatility and long term mean reversion.

Clearly, the scope for producing models in which agents are not fully rational is endless. In some behavioural finance models agents do not hold rational beliefs,

whilst in others they hold correct beliefs but make choices which are suboptimal in terms of subjective expected utility theory. The literature has indeed been criticised for a tendency to develop models that will explain observed phenomena by construction, rather than to develop models based on the theory and then test that against the data. The two models mentioned above make very different assumptions about investors' decision making and are able to explain some, but certainly not all of the anomalies in the data.

What all behavioural models have in common is that there is some additional source of variance in the return from holding equities beyond the variance in the future profit stream. This additional source of risk is an inefficiency in the market often known as noise trader risk (see De Long et al (1990 [32]) for the seminal account of noise trader risk).

In this chapter I set out to test whether a particular behavioural model can account for the time series data we have for the FTSE all-share index. The model that I use is a development of the behavioural model used by De Grauwe and Grimaldi (2006 [31]) to model the exchange rate. I choose this model for two reasons. Firstly, it models noise trader risk explicitly, using simple and general forecasting rules. Secondly, De Grauwe and Grimaldi use the model to explain the entire time series properties of an asset market rather than particular anomalies

(though they consider the exchange rate market, and only make general comparisons to the dynamics of real world markets, rather than providing a full test).

In this model, heterogeneous agents make a portfolio choice in order to maximise their utility. However, the expectations on which they base their choice are not rational in the conventional sense, but based on one of two simple heuristical rules. Agents choose to base their expectations either on a fundamental model of the asset price, or on a technical (or chartist) analysis of past asset price movements. Their choice of which rule to apply depends upon the past profitability of the rules. In this way the model can be viewed as being evolutionarily rational.

De Grauwe and Grimaldi (2006 [31]) have already established that the tension between simple rules can produce the type of dynamics that are characteristic of many asset price series in the real world: fat tails, excess kurtosis and GARCH properties. In this chapter I build on those results. I define the fundamental value of the FTSE as discounted future profits and am thereby able to discipline the model by requiring it to be consistent with the historical UK profits data. By bootstrapping the model, I am then able to determine whether it can account for the actual observed FTSE time series.

The question of whether asset market prices efficiently reflect fundamental values or whether they are partly the product of ‘investor sentiment’ is of key empirical importance. If the hypothesis that asset prices accurately reflect fundamental risks

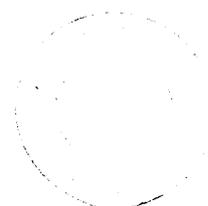
holds true, then all profitable speculation must be welfare enhancing and unfettered capital markets inevitably lead to the most efficient allocation of resources. If, on the other hand, the contentions of behavioural economists are correct, then speculation itself is a cause of risk - traders produce noise which is not present in the fundamental. In this way, capital market liberalisation is likely to give rise to new forms of risk. If the behavioural hypothesis is true, then this may well provide justification for some level of intervention in or regulation of capital markets.

The model is described in section 3.1. In section 3.2, I analyse the dynamics of the model and also examine the sensitivity of the model to its numerous parameters.

The model has a number of exogenous, behavioural parameters. In section 3.3 I examine the ways in which we can discipline our model by tying down these parameters empirically.

The data is presented in section 3.4. I examine the distribution of FTSE All-Share Index returns over a forty four year period. I also consider the fundamental driver of these returns, the UK profits series, over the same period.

In section 3.5 I stochastically simulate the model using the disciplining devices discussed, and test whether the model could reasonably explain the data observed in the real world. I also apply the same test to a benchmark rational model in order to compare my model's performance. Section 3.6 concludes.



3.1. A Behavioural Asset Pricing Model

In this section, I outline a behavioural model similar to that of De Grauwe and Grimaldi (2006 [31]). Heterogeneous agents make a utility maximising portfolio decision, conditioned upon their expectations of future asset prices. Their expectations of the future asset price is based on one of two simple heuristical rules - a fundamentalist rule or a chartist rule.

There are two assets in the model, a risky asset and a riskless asset. The riskless asset pays a known rate of return, r , per period. The risky asset's return is simply any appreciation in its price level, which is determined by market clearing.

We assume a continuum of agents from 0 to 1, each of whom is characterised by the forecasting rule, i , that they use. Each agent maximises their utility, which is a mean-variance function of wealth:

$$(3.1) \quad U(W_{i,t+1}) = E_{i,t}(W_{i,t+1}) - \frac{1}{2}\mu V_{i,t}(W_{i,t+1})$$

subject to the wealth constraint:

$$(3.2) \quad W_{i,t+1} = P_{t+1}d_{i,t} + (1+r)(W_{i,t} - P_t d_{i,t})$$

where μ is the coefficient of risk aversion, $d_{i,t}$ is the quantity of the risky asset held from period t until $t+1$, and $(W_{i,t} - P_t d_{i,t})$ is therefore the holdings of the risk free asset.

Given that $E_{i,t}$ is the expectation formed using forecasting rule i and conditioned on information available of time t (i.e. $d_{i,t}$, $W_{i,t}$, P_t and r are known with certainty); and that $V_{i,t}$ is a measure of one period ahead risk; then we can restate the problem as the maximisation of:

$$(3.3) \quad U(W_{i,t+1}) = d_{i,t}E_{i,t}(P_{t+1}) + (1+r)(W_{i,t} - P_t d_{i,t}) - d_{i,t}^2 \frac{1}{2} \mu V_{i,t}(P_{t+1})$$

The first order condition of this maximisation is:

$$(3.4) \quad \frac{\partial U(W_{i,t+1})}{\partial d_{i,t}} = E_{i,t}(P_{t+1}) - (1+r)P_t - d_{i,t} \mu V_{i,t}(P_{t+1}) = 0$$

Hence, the optimal holding of the risky asset by any agent is equal to its risk adjusted expected excess return:

$$(3.5) \quad d_{i,t} = \frac{E_{i,t}(P_{t+1}) - (1+r)P_t}{\mu V_{i,t}(P_{t+1})}$$

Aggregating across all agent types we get the market demand function:

$$(3.6) \quad D_t = \sum_{i=1}^I w_{i,t} d_{i,t}$$

$$(3.7) \quad = \sum_{i=1}^I \frac{w_{i,t} E_{i,t}(P_{t+1})}{\mu V_{i,t}(P_{t+1})} - (1+r)P_t \sum_{i=1}^I \frac{w_{i,t}}{\mu V_{i,t}(P_{t+1})}$$

where $w_{i,t}$ is the proportion of people who use forecasting rule i in time period t , and I is the total number of different forecasting rules used.

Market clearing implies that the total market demand for the asset must equal the supply of the asset, S_t . This allows me to solve for the market clearing price in period t :

$$(3.8) \quad P_t = \frac{-S_t + \sum_{i=1}^I \frac{w_{i,t} E_{i,t}(P_{t+1})}{\mu V_{i,t}(P_{t+1})}}{(1+r) \sum_{i=1}^I \frac{w_{i,t}}{\mu V_{i,t}(P_{t+1})}} + \eta_t \quad \eta_t \sim iid(0, \sigma_\eta)$$

where η_t is a white noise pricing error.

I must now describe the forecasting rules that individuals make use of. The fundamentalist rule forecasts that the market price will move towards its fundamental value, P^* , during the next period, unless it is already close to the fundamental, defined by the bounds $\pm C$:

$$(3.9) \quad \begin{aligned} E_{f,t}(P_{t+1}) &= P_{t-1} - \psi(P_{t-1} - P_{t-1}^*) && \text{where } |P_{t-1} - P_{t-1}^*| > C \\ &= P_{t-1} && \text{where } |P_{t-1} - P_{t-1}^*| \leq C \end{aligned}$$

We can think of C as an uncertainty bound. Fundamentalists recognise the uncertainty inherent in their modelling of the fundamental price and so they only take an active trading position when the actual price is significantly different from the fundamental value. If the actual price is already close to fundamental (i.e. within the band defined by C) then they predict no change in the actual price.

The chartist rule, on the other hand, simply extrapolates the previous movements of the market rate:

$$(3.10) \quad E_{c,t}(P_{t+1}) = P_{t-1} + \beta \sum_{j=1}^{\infty} \rho^{j-1} (1 - \rho) \Delta P_{t-j} \quad 0 < \rho < 1$$

Clearly, neither of these expectations forming processes is rational in the conventional sense. The contention of behavioural economics is that the level of complexity in the real world makes it impossible for agents to fully comprehend the markets in which they trade. In such a world, the ex-ante use of simple rules such as those in this model may constitute a best response. However, even in a complex world, the ex-post assessment of trading rules is relatively cheap. I therefore impose some limited rationality in the form of an evolutionary switching procedure based on the ex-post profitability of the competing rules. Agents are assumed to assess the ex-post risk adjusted profitability, $\Pi'_{i,t}$, of each of the forecasting rules and then select the rule that they will use in the next period. Hence, the proportions of agents using each of the rules develops according to the following identities:

$$(3.11) \quad w_{f,t} = \frac{\exp(\gamma \Pi'_{f,t})}{\exp(\gamma \Pi'_{f,t}) + \exp(\gamma \Pi'_{c,t})}$$

$$(3.12) \quad w_{c,t} = \frac{\exp(\gamma \Pi'_{c,t})}{\exp(\gamma \Pi'_{f,t}) + \exp(\gamma \Pi'_{c,t})}$$

$$(3.13) \quad \Pi'_{i,t} = \Pi_{i,t} - \mu\sigma_{i,t}^2$$

$$(3.14) \quad \Pi_{i,t} = [P_{t-1} - (1+r)P_{t-2}].\text{sign}[E_{i,t-1}(P_t) - (1+r)P_{t-1}]$$

$$(3.15) \quad \sigma_{i,t}^2 = \sum_{j=1}^{\infty} \rho^{j-1}(1-\rho) [E_{i,t-j-1}(P_{t-j}) - P_{t-j}]^2$$

where $w_{f,t}$ is the proportion of agents at time t using the fundamentalist rule and $w_{c,t}$ is the proportion using the chartist rule. γ is a parameter measuring the intensity of revision of the forecasting rules. If $\gamma = 0$ then agents never change the forecasting rule that they use, and exactly half the population uses each rule. As γ approaches infinity all agents switch immediately to the rule that was most profitable in the preceding period. For all intermediate values agents switch between rules, but only sluggishly. This suggests some form of status quo bias, as suggested by Tversky and Kahneman (1974 [91]).

The final equation I need to close the model is the determination of the fundamental price. When I come to test the model empirically, I will use the underlying profits process to discipline the fundamental price. However, in the following analysis of the dynamics of the model, it is enough to observe that the fundamental price of the asset, P_t^* , follows a random walk with drift:

$$(3.16) \quad P_t^* = P_{t-1}^* + \delta + \varepsilon_t \quad \varepsilon_t \sim iid(0, \sigma_\varepsilon)$$

I assume that the news, ε_t , becomes public information during the time period.

The timing of events within a period can be summarised as follows:

- (1) The beginning of period information set is given by $\Phi_t = \{P_{t-i}, P_{t-i}^*\}_{i=1 \text{ to } \infty}$.
- (2) The profitability and risk of the various forecasting rules is assessed, $\Pi_{i,t}$ and $\sigma_{i,t}^2$, and agents decide which forecasting rule to use this period. Hence, the weights $w_{i,t}$ are determined.
- (3) Expectations are formed of the next period price, $E_{i,t}(P_{t+1})$, and so this period's demand functions, $d_{i,t}$, are determined.
- (4) Market clearing determines the present period price, P_t .
- (5) News determines the present period fundamental, P_t^* .

3.1.1. Solving the Model

Solving the model is relatively straightforward. It is simply a matter of iterating the equations of the model, given a particular parameterisation and realisation of the shocks.

I use the Fortran programming language to simulate the model. The Fortran code is included in appendix A.1, but this code also includes the estimation and testing procedures that I will discuss in subsequent sections.

3.2. Basic Analysis of the Model

The highly non-linear nature of the model makes the analysis of its steady state and dynamics very complex. In this section I therefore take a three pronged approach to the analysis. Firstly, I look at an example stochastic simulation in order to illustrate some of the basic properties of the model. Secondly, I consider a simplified, ‘deterministic’ version of the model, so as to be able to provide some analytical results about the steady state. Thirdly, I carry out some numerically simulated sensitivity analyses of the full model in order to provide some intuition about the role of the model’s parameters and initial conditions.

3.2.1. An Example Stochastic Simulation

Figure 3.1 shows an example of a 1,000 period stochastic simulation of the model. The fundamental value of the asset, P_t^* , is given by the black line in the upper portion. As discussed earlier, this is a random walk with drift. The pink line shows the market price for the asset, P_t . In this simulation the asset price shadows its fundamental value less than precisely for the majority of the simulation period, but at times it deviates markedly and persistently away from the fundamental. The lower portion of figure 3.1 shows us the weight on the fundamentalist rule, w_f , over the simulation period. We can see that the deviations from fundamental

value (or bubbles) are associated with low weight on the fundamentalist rule, and hence high weight on the chartist rule.

3.2.2. A Simple, ‘Deterministic’ Model

I simplify the model in a number of ways so as to provide an analytical solution:

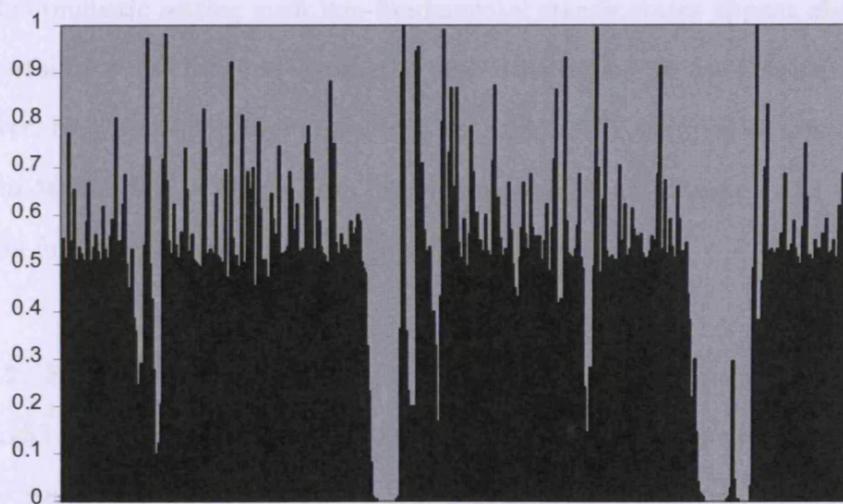
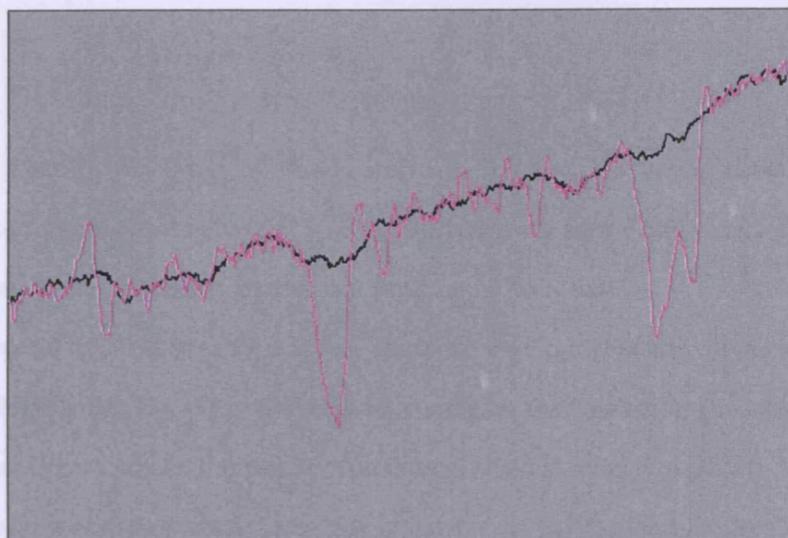
- I set the net supply of the asset equal to zero, $S_t = 0$.
- there is no pricing error, $\eta_t = 0$.
- there is no news in the fundamental price, which I normalise to zero, $P_t^* = 0$. In this way, P_t can be considered as the discrepancy between the actual price and its fundamental value.
- the interest rate is equal to zero, $r = 0$.

I can then express the model as a difference equation:

$$(3.17) \quad P_t = [1 - \psi\Theta_{f,t} + \beta(1 - \rho)\Theta_{c,t}] P_{t-1} - [\beta(1 - \rho)^2\Theta_{c,t}] \sum_{i=2}^{\infty} \rho^{i-2} P_{t-i}$$

where $\Theta_{i,t} \equiv \left[\frac{w_{i,t}/\sigma_{i,t}^2}{w_{f,t}/\sigma_{f,t}^2 + w_{c,t}/\sigma_{c,t}^2} \right]$ are the risk-adjusted weights of the different trading rules [NB $\Theta_{f,t} + \Theta_{c,t} = 1$].

Figure 3.1. An example stochastic simulation



3.2.2.1. Steady State Solutions. By dropping the time prefixes, I can solve for all of the model's steady state price levels, \bar{P} :

$$(3.18) \quad \bar{P}\psi\Theta_f = 0$$

I can now identify two types of steady state. In the first, the actual asset price is equal to its fundamental value ($\bar{P} = P^* = 0$); there are zero profits; zero risk; and the weights on the forecasting rules are both equal to a half.

The second type of steady state is characterised by a weight of zero on the fundamentalist rule ($\Theta_f = 0$). A sufficient condition for this to be the case is that chartist risk is zero whilst the risk for fundamentalists is strictly positive. This can occur with any constant asset price not equal to the fundamental value. Clearly, in a deterministic setting such non-fundamental steady states appear absurd - it is uncertainty caused by the stochastics that drive agents to use heuristical rules. However, by considering the model with the stochastics stripped out, we can get some understanding of how shocks can lead to a persistent divergence of the price from its fundamental value.

3.2.2.2. Sensitivity Analyses. In this section I consider how the parameters in the model affect the way in which the asset price responds to a single shock.

We can see clearly from equation 3.17 that given an initial price of zero, future prices will also be zero. I therefore shock the initial period price in order to see

the effect on the steady state price of different parameterisations of the model (in this case I set $P_0 = 5$ and $P_t = 0$ for all $t < 0$). In each case I begin with a base parameterisation and then vary one of the parameters at a time in order to analyse the sensitivity of the steady state solution. Figure 3.2 shows how the steady state price changes as we vary each parameter. The red dotted lines represent the base parameterisation of $\psi=0.3$, $\beta=0.85$, $\rho=0$, $\gamma=6.0$, $\mu=1.0$ and $C=0$, under which the steady state price is 14.2.

In order to understand the relationships shown we need to appreciate the self-fulfilling nature of the model. If agents predict a large increase in price, then they will demand more of the asset and this will, in turn, drive the price upwards. Hence, when ψ is high, representing an aggressive fundamentalist forecast, the price is driven back to its fundamental value following the shock. On the other hand, when β is high and chartists expect a high correlation between future and past price changes, the price diverges from fundamental.

Another important characteristic of the model is the herding behaviour encapsulated in the switching rule. This implies the existence of a tipping point in the model. If the risk-adjusted profits associated with the chartist rule are just marginally higher than those associated with the fundamentalist rule then agents will begin to shift towards the chartist rule and the chartist forecast will become self-fulfilling. If, on the other hand, the fundamentalist rule proves more profitable, agents will flock towards it and the price will move towards its fundamental value. In this way, the relative profitability of the two forecasts in the period immediately

following a shock is key. The tipping point is illustrated in the sensitivity analyses in the discontinuities in the relationships between each of the parameters ψ and β and the steady state price.

ρ represents the weight that chartists put on previous price changes relative to last period's price change. Therefore, as ρ increases the reaction of the chartist forecast to any individual price shock is increasingly spread out over time. As ρ increases and the reaction is staggered out over more periods, there are two effects:

(1) there is a less aggressive reaction by individual chartists in the period immediately following the shock.

(2) there is a more aggressive reaction to the shock in subsequent periods.

This is significant because, when the chartist rule is the more profitable in the period immediately following the shock, many agents move to the chartist rule in subsequent periods, and so the staggered reaction to the shock is acted upon by a greater number of agents. Hence, the chartist rule is more vigorously reinforced.

In our sensitivity analysis the second effect dominates for values up to $\rho = 0.51$, and the steady state price gets further away from the fundamental as ρ increases. Beyond this point, the first effect dominates and the steady state price gets closer to the fundamental until the model reaches the tipping point discussed above, at $\rho = 0.62$.

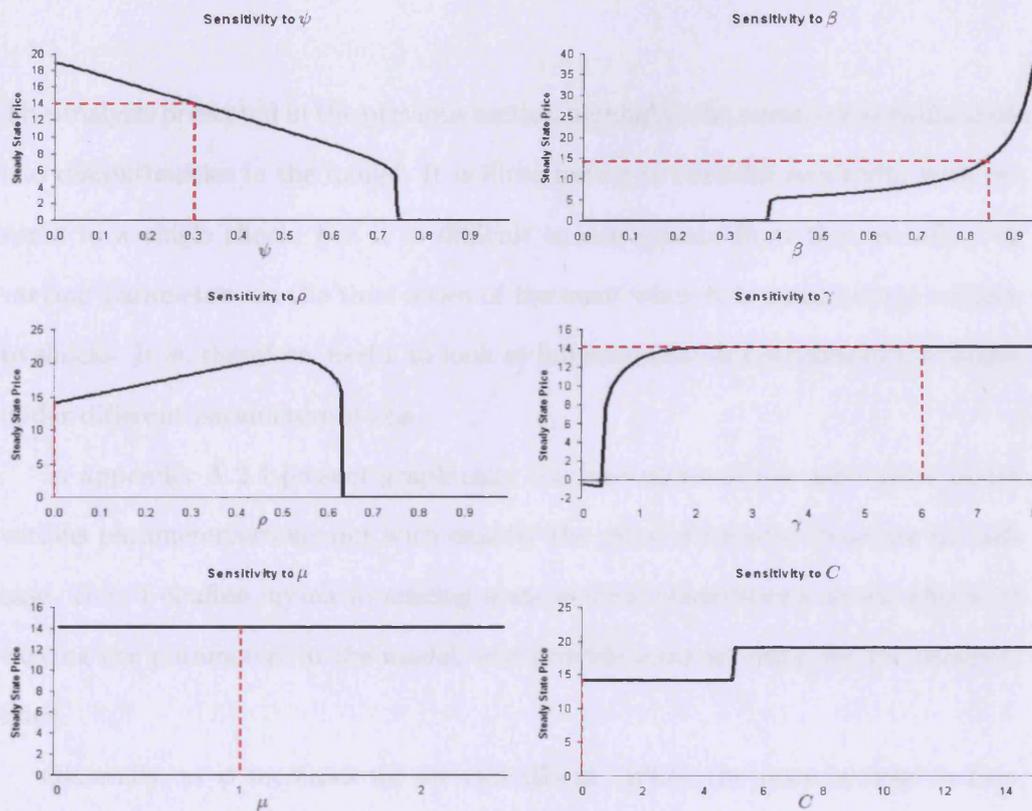
When $\gamma = 0$ there is no switching at all, so 50% of agents always use each of the forecasting rules. In the present case, the chartists' expectations cause the price to increase initially, beyond the magnitude of the shock. The presence of agents using the fundamental rule means that the price is pulled back towards its fundamental value once the effect of the shock dies out in the chartist forecast. However, the chartist forecast causes the price to overshoot the fundamental. The price finally comes to rest at a value just below fundamental. At this price, fundamentalists expect an appreciation in the asset price but also experience a positive risk in their forecast, whilst chartists' expectations are fulfilled. For each agent, demand is constant and, given that the price is as well, profits are zero for all agents.

As γ increases agents begin to switch away from the less profitable rule. In the period immediately following the price shock, the chartist rule is more profitable and so agents switch to it. When γ is positive but small this switching occurs very slowly and so serves only to increase slightly the overshooting effect. When $\gamma > 0.34$ the switching occurs rapidly enough that the pull back towards the fundamental is short-lived, and the initial chartist expectations become self-fulfilling.

The coefficient of risk aversion, μ , has little role to play in this 'deterministic' version of the model and so the sensitivity analysis is flat.

As long as C is less than the size of the shock, fundamentalists behave in the same way as if $C = 0$. However, for C greater than the size of the shock fundamentalists exert no influence on the price during the period immediately following the shock. Hence, we have a step shaped sensitivity analysis.

Figure 3.2. Sensitivity analysis in the deterministic model



One of the striking features of all these graphs are the discontinuities in the relationships between the parameters and the steady state outcomes. It is these discontinuities that cause chaotic behaviour in the model - very small causes can have very large effects.

3.2.3. Further Stochastic Simulation

The analysis presented in the previous section highlights the numerous non-linearities and discontinuities in the model. It is illuminating to consider sensitivity with respect to a single shock, but it is difficult to extrapolate from this the effect of varying parameters on the time series of the asset when it is continuously subject to shocks. It is, therefore, useful to look at full stochastic simulations of the model under different parameterisations.

In appendix A.2 I present graphically the time series of the asset price under various parameterisations but with exactly the same stochastic structure in each case. Here I confine myself to making some general observations on the effects of varying the parameters in the model, and provide some intuition for the observations.

Generally, as ψ increases we see two effects. When the price is close to fundamental, it follows the fundamental value more closely because of the increased aggressiveness in the fundamentalist forecast. However, when bubbles occur they are more persistent and deviate further from the fundamental because once agents shift to the chartist rule they are less likely to shift back to an aggressive fundamentalist forecast.

For low values of β the asset price closely shadows its fundamental value. As β increases bubbles becomes more common, more persistent and more extreme.

Increases in the value of ρ ameliorate the response of the chartist rule to individual shocks and so they reduce the magnitude of the divergence of the asset price from fundamental during bubbles.

For high values of γ agents switch rapidly between the forecasting rules in response to relative profitability. This makes the asset price more volatile, and divergences from fundamental more extreme.

As risk aversion, μ , increases we see more extreme and more persistent deviations from fundamental. Given that chartist behaviour introduces additional risk into the market, in the form of noise, it appears to be counter-intuitive that increased risk aversion encourages chartist behaviour. However, the chartist forecast covaries more closely with the price than does the fundamental rule. Therefore, as the coefficient of risk aversion increases, more agents will choose to use the chartist rule and this causes longer and more extreme bubbles.

When $C > 0$ fundamentalists exert no influence over the asset price so long as it is within $\pm C$ of the fundamental. However, so long as the price remains within those bounds fundamentalists are likely to be more or less equally profitable as the chartists. Hence, when the price hits the bound, there are likely to be around 50% fundamentalists in the market who will drive it back within the bounds. Hence as C increases we see more volatility in the price within the bounds but fewer occurrences of the price escaping the bounds.

3.3. Disciplining the Model

It is clear from the foregoing discussion that we have a very flexible model. With six exogenous parameters and two exogenous stochastic processes, we have a model with so many degrees of freedom that it would be surprising if we could not, under some parameterisation, provide a good fit for real world data. In this section I will discuss some empirical limitations on the exogenous aspects of the model.

3.3.1. The Fundamental Price

A crucial aspect of the model is the fundamental price on which fundamentalists base their expectations of future price movements. We can think of this as being akin to a rational expectation of the price level, based on the fundamental value of the asset. In other words, it is a discounted sum of the future cashflows arising from ownership of the asset. In the case of equities, this is equivalent to the discounted sum of future profits since all profits must ultimately accrue to the equity owners.

Clearly, profits are driven by the marginal productivity of capital. However, there is no generally accepted theory explaining the growth of the productivity of capital. Conventionally, therefore, both productivity and profits have been modelled by simple univariate time series processes. I follow this convention, and use

a parsimonious ARIMA representation of profits as the basis for the fundamental value of the asset.

In this way I can provide some discipline to the fundamental price in the model - it must be consistent with the actual profits process that underlies it. And, of course, we have data for this profits process. In section 3.5 I will derive the parsimonious univariate time series representation of the UK profits series. In the following section, I will then derive bootstrapped series of the fundamental price which are consistent with rational expectations of future profits.

3.3.2. The Coefficient of Risk Aversion

A significant amount of research has been carried out in respect of the empirical magnitude of the coefficient of risk aversion (e.g. Epstein and Zin 1991 [35]). The general consensus suggests a value of unity. Accordingly, I set μ equal to 1.

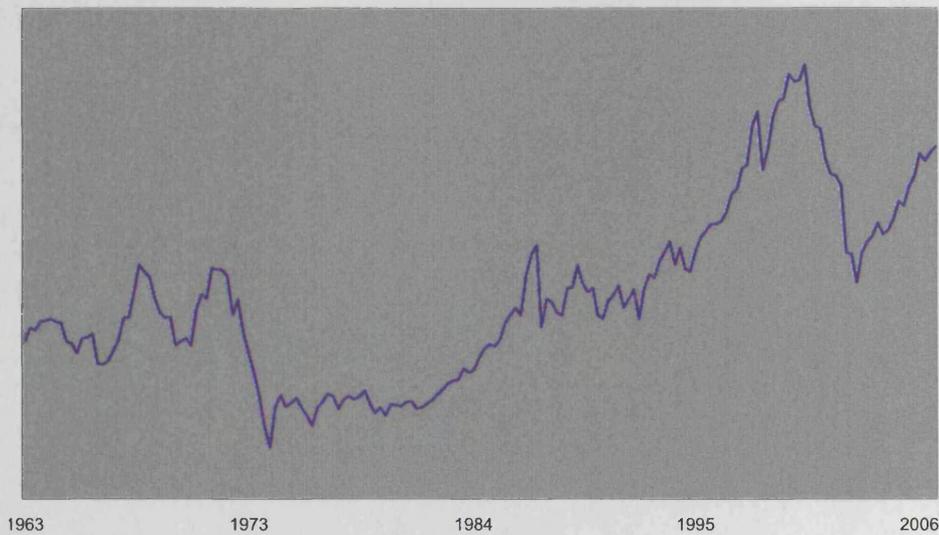
3.3.3. Other Parameters

The value of the parameters in the forecasting rules, ψ , β and ρ , and the switching rule, γ , are clearly unobservable and more open to debate. Their value will certainly depend upon the length of time period we consider, and must fulfil some concept of 'reasonableness'. For example, when I come to test the model I will consider

quarterly time periods because profits data are only available quarterly. ψ then measures the proportion of the gap between market price and fundamental value that fundamentalists believe will be closed during the next three months. We might then say that values of ψ less than 0.1 are unreasonable. Beyond this, all we can say is that the behavioural parameters must be internally consistent and ensure stability in the model i.e. $0 < \psi < 1$, $0 < \beta < 1$, $0 \leq \rho < 1$ and $0 \leq \gamma$.

As far as the uncertainty parameter C is concerned, it is difficult to define what would be a reasonable size band within which fundamentalists remain inert. As a reasonable starting point, I restrict C to being no more than 10%.

Figure 3.3. Real FTSE 1963-2006



3.4. The Data

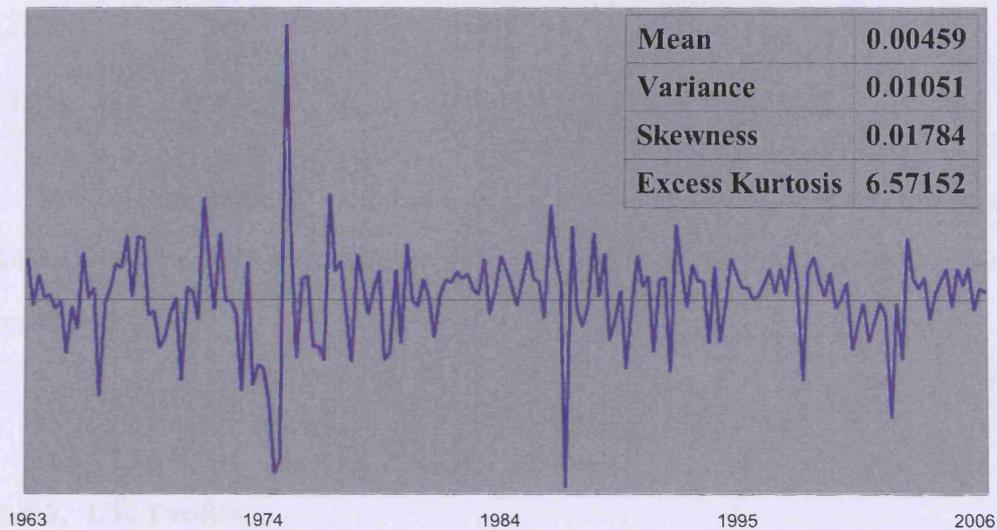
3.4.1. The FTSE All Share Index

The FTSE All-Share Index (hereon referred to as the FTSE) is considered to be the best performance measure of UK equities, accounting for 98% of the UK's market capitalisation. The index was launched in April 1962. Figure 3.3 illustrates the real value of the index (calculated using the GDP deflator) from the first quarter of 1963 up until the last quarter of 2006.

I approximate the real rate of return to the FTSE as the first difference in logs:

$$(3.19) \quad FTSE \text{ rate of return} = \Delta(\ln[\text{real } FTSE])$$

Figure 3.4. FTSE real quarterly rates of return and estimated population moments



This series, along with its estimated population moments, are shown in figure 3.4. It is clear from the illustration that these returns are not normally distributed. The sample is characterised by fat tails, particularly of negative returns, positive skewness and excess kurtosis.

I carried out a regression analysis to determine the best fit time series for the rates of return. Using the Hannan-Rissanen procedure with Schwartz selection criterion, I found that the following ARMA(0,0)-GARCH(1,1) process provided the most parsimonious representation of the series:

$$(3.20) \quad \Delta(\ln[\text{real } FTSE_t]) \underset{\text{(std error)}}{=} \underset{(0.0064)}{0.0080} + u_t$$

$$(3.21) \quad \sigma_{u_t}^2 = \underset{(0.0011)}{0.0016} + \underset{(0.0795)}{0.1983}u_{t-1}^2 + \underset{(0.1774)}{0.6343}\sigma_{u_{t-1}}^2$$

Hence I can conclude, as is usual in the finance literature, that stock returns follow a random walk with drift but that the volatility of returns is to some extent predictable.⁵

3.4.2. UK Profits

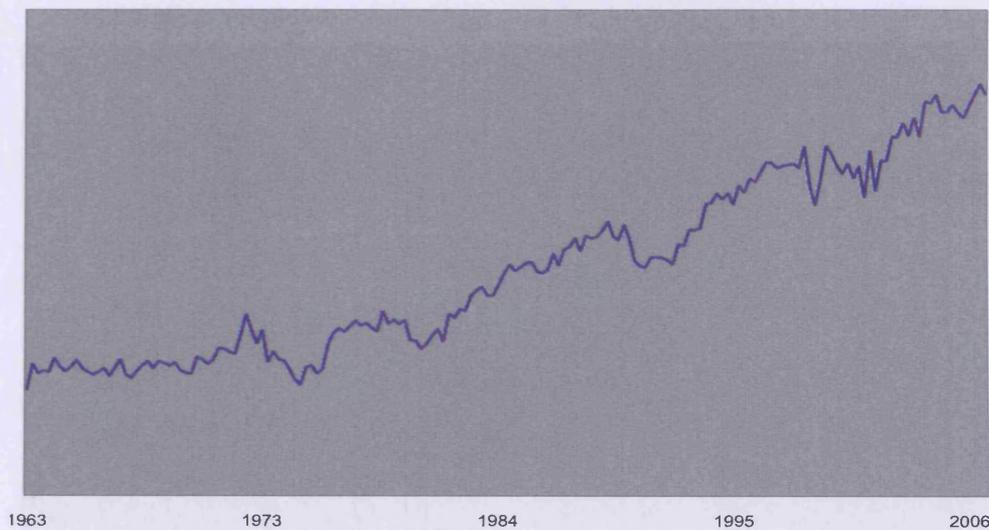
The fundamental variable underlying the performance of the FTSE is the profits of UK publicly quoted companies. We use a seasonally adjusted transformation of the Office for National Statistics' quarterly profits series⁶ as a measure. Figure 3.5 shows this series in real terms (calculated using the GDP deflator).

In order to investigate the time series properties of the data I take the differences of the log of profits to produce a stationary series representing the rate of

⁵Despite the fact that the constant terms in both the primary and the variance equations are insignificant, we report them here because when we come to test our behavioural model we will want to see whether it is capable of reproducing such values as have actually occurred in the data.

⁶The series that we use is the sum of the gross operating surplus of private non-financial corporations (ONS code NRJK) and the gross operating surplus of financial corporations (ONS code NQNV). Both series are reported in table 1.3 of the Monthly Digest of Statistics and on the ONS website. We use a multiplicative moving average method to adjust for the seasonality in the raw data.

Figure 3.5. Real UK Profits 1963-2006

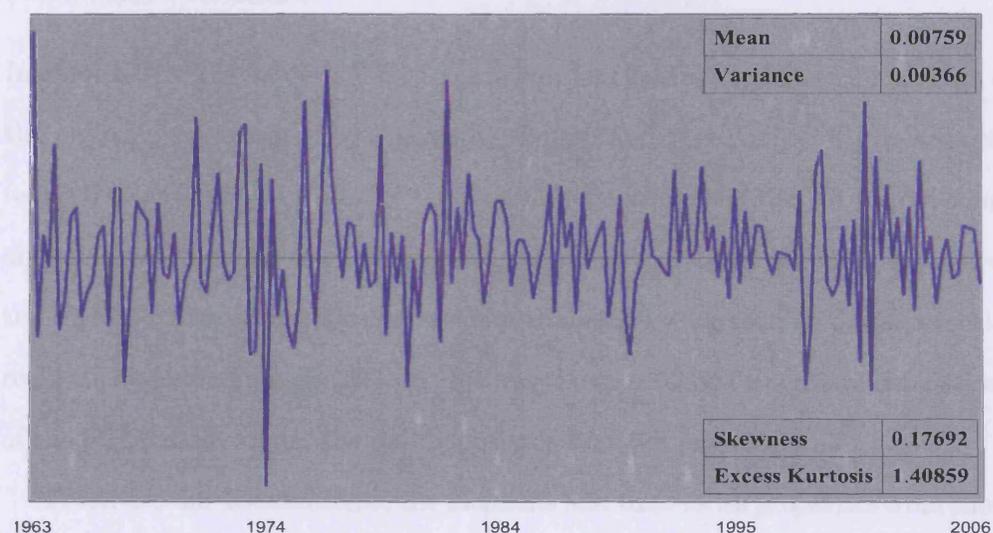


change of quarterly profits. The transformed series and the related estimates of the population moments are shown in figure 3.6.

I once again used the Hannan-Rissanen procedure with Schwartz selection criterion to determine the most parsimonious ARIMA representation of this series. As is clear from figures 3.5 and 3.6, the log profits series is integrated of order 1. The best representation of the rate of change of profits is given by the following ARMA(1,0):

$$(3.22) \quad \Delta(\ln[\text{real profit}]_t) = \underset{\text{(std error)}}{0.008435} - \underset{\text{(0.004325)}}{0.260975} \Delta(\ln[\text{real profit}]_{t-1}) + u_t$$

Figure 3.6. Rate of change of real quarterly profits and estimated population moments



Diagnostic testing shows that the residuals from this regression do not display significant low order serial correlation (Breusch-Godfrey test), nor do they display ARCH properties. However, we can reject the hypothesis that the residuals are normally distributed (Jarque-Bera p -value of 0.034). The residuals have positive skewness of 0.08 and kurtosis of 3.95. This non-normality implies that, when we come to simulating the profits process in our model-testing procedure, we need to bootstrap these residuals rather than use stochastic simulation.

3.5. Testing the Model

3.5.1. The Procedure

In order to test the model's ability to account for the empirical facts, I make use of the methodology adopted by Meenagh, Minford and Peel (2007 [68]) (henceforth referred to as MMP). I begin by taking 50,000 bootstraps of the UK profits series and discount them to provide possible realisations of the fundamental value of the FTSE. I then stochastically simulate the model with each of these potential realisations of the fundamental. In this way, I derive 50,000 stochastic simulations of the FTSE series under the null hypothesis that the model is true.

I then use the distribution of the moments and time series properties from these simulations to construct 95% confidence intervals. If I find that the moments and time series properties of the actual FTSE series lie outside the confidence intervals then I can reject the null hypothesis that the model is true. Conversely, if the properties of the actual FTSE lie within the confidence bounds, then I cannot reject the model. I also employ a joint test of all the moments and GARCH parameters.

In fact, I carry out this testing procedure both on the 'behavioural' model and on MMP's 'rational' model using the same dataset⁷ so that I can provide a comparison with a RE-EMH model.

⁷This dataset differs from that of MMP in two ways. Firstly, the MMP data covers 1963Q1 to 2002Q2. My data continues up to 2006Q4. Secondly, MMP use the average daily close over the quarter as their measure of the FTSE price whereas we use the quarterly close. Clearly, our FTSE data is therefore more volatile.

I use a search algorithm to find the best-fitting values for the behavioural parameters and the standard deviation of the pricing shock in the behavioural model (subject to the restrictions discussed in section 3.3). The parameterisation used is as follows⁸:

ψ	β	ρ	γ	μ	C	σ_η
0.425	0.99	0.2	3.75	1.0	0.5	0.025

3.5.2. The Results

Table 3.1 summarises the moments and time series properties of the FTSE simulations from each model and compares them to the actual data. As can be seen, the only case of a property of the actual data falling outside the confidence bounds is variance in the MMP model. For all other properties, the MMP model passes the 95% test. The behavioural model satisfies every property test. Graphical representations of the distributions of properties under each model are shown in appendix A.3.

According to the results in figure 3.1 we cannot reject the behavioural model, but can reject the MMP model on the basis that less than 5% of the simulations of that model produce variances in returns as large as we see in the actual FTSE index. It is clear, however, that the MMP model does better than the behavioural

⁸In the model, C is measured in absolute terms. An absolute value of 0.5 is equivalent to 10% of the steady state fundamental price.

Table 3.1. Moments and time series properties of the simulated and actual FTSE series

	ACTUAL	Behavioural Model		MMP Model	
	FTSE	Lower	Upper	Lower	Upper
		2.5% limit	2.5% limit	2.5% limit	2.5% limit
Mean	0.0046	-0.1403	0.1389	-0.0075	0.0139
Variance	0.0105	0.0009	0.0155	0.0026	0.0076
Skewness	0.0178	-1.3116	2.2961	-0.7626	1.5572
Kurtosis	9.5715	1.9121	14.1262	2.6187	18.8049
Trend	0.0080	-0.1517	0.1459	-0.0079	0.0153
ARCH constant	0.0016	0.0000	0.0016	0.0001	0.0054
ARCH	0.1983	0.0000	0.9956	0.0000	0.5516
GARCH	0.6343	0.0000	0.9419	0.0000	0.9232

model in terms of some of the other properties. For example, if we look at the distribution of the mean return in each model, the mean of the distribution is close to the actual FTSE mean return in each case, but the distribution is much tighter around that value in the case of the MMP model. This observation suggests two things. Firstly, it would be useful to consider a single metric that takes into account the joint distribution of all of the properties considered in figure 3.1 and see whether the FTSE lies within the 95% confidence interval for such a metric. We could also use the p-value of the metric to compare our models. Secondly, it would be useful to be able to say something about the power of our tests.

It is important to note at this point that the MMP model involves a latent profits model. The profits process is driven by the model parameters via a Markov switching process. Any joint test of this model should, strictly speaking, consider the time series properties of the profits process as well as those of the asset price.

In the behavioural model, on the other hand, the profits process is simply driven by the historical profits data, and so there is a good fit by construction. In order to compare the two models, here we only consider a test of the joint distribution of the asset price characteristics. This could prove to be a lenient treatment of the MMP model, but we simply accept that the parameters of that model were selected in the first place to fit the profits data as well as the asset price data.

Minford, Theodoridis and Meenagh (2007 [70]) propose such a metric, based on the Mahalanobis distance. This is a statistical distance measure based on the correlations between different variables. Formally, the Mahalanobis distance is calculated as $MD_i = \sqrt{(x_i - \mu)Cov^{-1}(x_i - \mu)'}$, where x_i is a vector of observations (in this case, the vector of moments and time series properties from simulation i), μ is the mean of the x_i , and Cov is the associated covariance matrix. I calculate the MD of each simulation from the set of model simulations and use the distribution of those values to get a 95% confidence limit. I then calculate the distance of the actual data from the model simulations, and see whether it lies within the bound. For each model, the actual FTSE lies within the 95% limit (see appendix A.3 for a graphical representation).

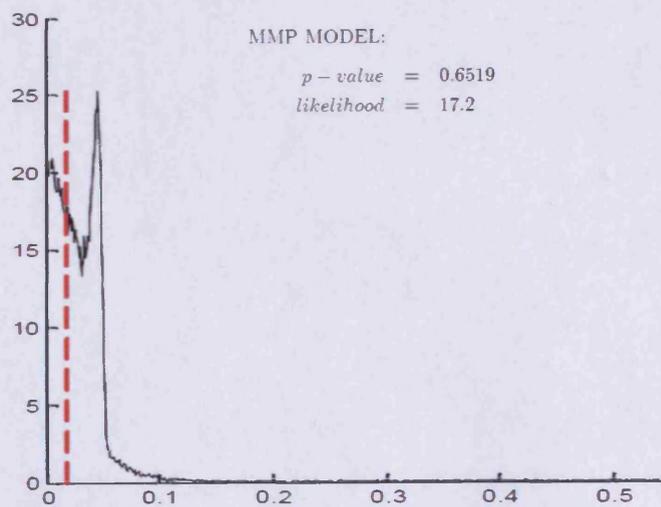
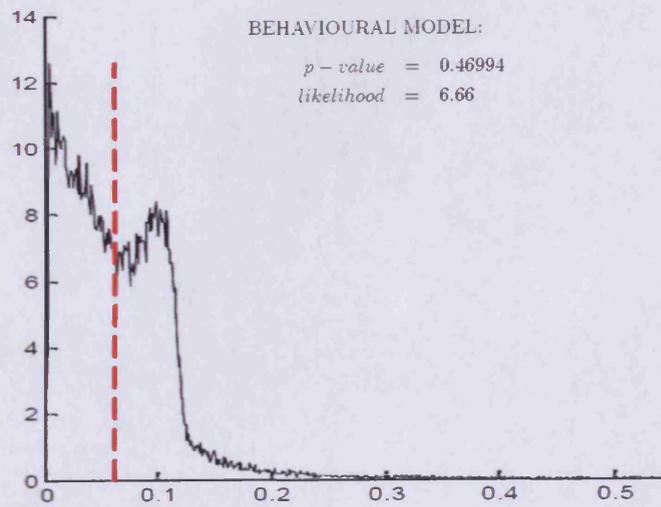
If I can find a p-value for the distribution of Mahalanobis distances then I can use this to rank the alternative models. If I can find a likelihood value then I can use this to give an idea of the power of the test. It is a well known result that if the population under consideration may be considered multivariate normal, then the square of the Mahalanobis distance will follow a chi-square distribution.

However, as can be seen in appendix A.3, the distributions of many of the time series properties are highly non-normal. For example, for each model the GARCH parameter is insignificant in 30% to 55% of the simulations. Unsurprisingly, I find that the chi square distribution is a very poor fit for our distributions of MD . I therefore simply use the actual distribution of MD to calculate p-values and likelihoods for each model. I normalise the MD distribution so that it is bounded by 0 and 1 and so that the area under the distribution is equal to one. The normalised distribution for each model is shown in figure 3.7. The dashed red line in each graph shows the MD of the actual FTSE, and p-values and likelihoods are as stated.

To summarise the various tests, I can say that the MMP model has more power than the behavioural models we look at (illustrated by the concentration of MD values and the narrowness of the confidence intervals for individual time series properties). It also has a reasonably good fit overall, with a p-value of 65%. However, it fails to account for the high volatility of FTSE returns that we see in the data.

The behavioural model also provides a good fit overall, with a p-value of 47%, and can account for all of the individual time series characteristics of the data. We should note, however, that the test appears to have less power when applied to the behavioural model.

Figure 3.7. p-values and likelihood values for each model



In conclusion, it is difficult to distinguish between the models on the basis of their overall fit of the data. However, it is clear that the MMP model cannot account for the large variance in returns, whereas the behavioural model can.

3.6. Conclusions

In this chapter I have provided a framework with which to test the efficacy of different asset pricing models in explaining the empirical facts. I have applied this framework to two competing models; a behavioural model which has been discussed in detail, and a RE-EMH model which has been explained elsewhere (Meenagh, Minford and Peel 2007 [68]). The results show that either model gives a reasonable account of the overall FTSE time series, but that the MMP model struggles to explain the volatility in stock returns.

Essentially the choice between the two modelling frameworks comes down to whether the complex dynamics of asset prices are driven by a complex profits process or by complex behaviour in asset markets. The implications of the two models are starkly different. If we accept the first argument then asset prices reflect the true riskiness of asset ownership. If we accept the second argument then there is a source of inefficiency in equity markets, in the form of noise trader risk.

The behavioural model explains the dynamics of the asset price as the result of the tension between fundamental and chartist forecasting rules. There is no doubt that the model put forward in this paper is far from a precise account of the way that asset markets function. There is a wealth of evidence (e.g. Brock et al 1992 [24]) that technical or chartist analysis has a prominent role on the trading floor. It is, however, clear that the chartist rules employed by traders are far more

sophisticated than that which I put forward in this chapter (equation 3.10). There is clearly scope for developing more sophisticated behavioural models, and they may well provide a better fit for the data.

What is significant to note is that even with the extremely simple forecasting rules that I use here, we have a model that produces very complex dynamics which are consistent with the actual data. Clearly, what leads us to reject this model as a precise representation of the actual FTSE market is the intuition that such highly rewarded financial practitioners could not be using such naive forecasting rules. A useful extension to this paper would be to examine more closely the rules that practitioners do actually appear to use in the real world (e.g. volatility modelling) and incorporate them into our framework.

A concern that we may have about the results in this chapter is that the power of the test appears to be low. Of course, it is the chaotic nature of the model that gives it this property, and this is exactly what I am exploiting in order to explain the FTSE dynamics.

A more particular concern is that the model allows for extremely long bubbles, where the asset price deviates massively from its fundamental value. This occurs when all of the agents in the model switch to a chartist rule. In this case, the model's stochastics provide the only route back towards the fundamental value. This is counter intuitive. We may believe that the market can be dominated by non-fundamental beliefs in the short run but it is common to think that there is some mechanism that ensures a return to fundamental value in the long run. The

rational expectations literature makes use of transversality conditions to eliminate the possibility of long-run bubbles caused by self-fulfilling expectations. It may be that we need something similar here to prevent long-run self-fulfilling chartist bubbles.

Analysis of the model shows that the parameter C is very important in the explanation of the FTSE dynamics, in particular the high kurtosis. This is because, with positive C , fundamentalists don't expect any movement back towards the fundamental value if the asset price is already close to it. They only expect a return towards fundamental if there is already a large discrepancy between the actual and the fundamental values. I have explained this as a consequence of the uncertainty surrounding the fundamental value. An equivalent mechanism might be found in the tension between short run and long run behaviour. If agents choose to follow the market in the short run (following chartist rules) but form forward looking expectations in the long run (following fundamental rules), this may have a similar effect to ignoring small discrepancies from the fundamental but acting upon large discrepancies. Such behaviour could potentially be accounted for by the relative infrequency of news about the fundamental as compared to the high frequency of news about the market, and the resulting effects on behaviour, particularly in the presence of principal-agent problems.

Incorporating forecasting rules that model this short run versus long run dichotomy would be a worthwhile extension to the present model.

The MMP model explains the dynamics of the asset price simply as the discounted sum of future profits, but where those future profits are driven by a Markov switching procedure with four regimes. The fourth regime is a massive crash with very low probability that has not been realised in the observable data. Extensions to this framework are harder to envisage.

Given that I find that this simple behavioural model performs at least as well as the MMP model, and given that the scope for extending and improving the behavioural model is great, I conclude that behavioural rules have a useful role in explaining the dynamics of asset prices. I would further argue that this is particularly true in the short run, when news about fundamentals is infrequent. This conclusion implies that there is some source of inefficiency in equity markets, namely noise trader risk. The question then arises as to whether there are policies that could be put in place to mitigate this inefficiency. This question provides the motivation for the next chapter.

CHAPTER 4

The Monetary Policy Implications of Behavioural Asset Bubbles

In this chapter my intention is to examine the implications that behavioural finance theories may have for the implementation of monetary policy. In particular, I hope to inform the debate on whether and to what extent asset prices should be included in central banks' policy considerations.

The use of asset prices in monetary policy formulation has a long and chequered history. For the majority of history, monetary policy has been inextricably linked to asset prices. The value of money was tied either to the value of precious metals or to the value of other currencies almost continuously up until the final decades of the twentieth century.

However, a near-consensus evolved within the literature over the latter decades of the twentieth century. It held that monetary policy should respond to expected inflation and possibly to the output gap, but should not be directly influenced by asset price movements. There has been some dissent to this view (for example Cecchetti et al (2000) [27]) and, furthermore, some evidence exists to suggest that

central banks do indeed take asset price movements into consideration when setting interest rates (Mishkin (2007) [71], Cecchetti et al (2000) [27]).

I will begin with a review of the literature that reflects both sides of this debate. I will then devote most of the chapter to developing a dynamic stochastic general equilibrium (DSGE) model based on heuristical foundations of the type that I considered in the previous chapter. I will use the results from the model to address the issue of whether the absence of rational expectations in financial markets may provide a justification for including asset prices in central banks' Taylor rules¹.

¹I use the phrase Taylor rule throughout this chapter to refer generically to monetary policy rules.

4.1. Literature Review

In his now infamous speech to the American Enterprise Institute, the then Federal Reserve chairman Alan Greenspan posed a question very similar to the one I am now attempting to answer:

But how do we know when irrational exuberance has unduly escalated asset values, which then become subject to unexpected and prolonged contractions as they have in Japan over the past decade? And how do we factor that assessment into monetary policy? (Greenspan 1995 [48])

These comments were, at the time, a rare reflection on the potential importance of asset mis-pricing to monetary policy. The slumps in world stock markets which have followed, in the early 2000s and during the recent credit crunch, have strengthened the view that asset price bubbles exist as an empirical fact. However, the literature relating to the subject is still relatively sparse.

Bernanke and Gertler (2000 [16]) provide the original investigation into the implications of asset price bubbles for monetary policy. They incorporate an exogenous asset bubble into their financial accelerator model (Bernanke, Gertler and Gilchrist 1999 [18]). The financial accelerator model is a New Keynesian model in which firms own their stock of capital directly, rather than renting it from households as is usual in most models. It also differs from the standard New Keynesian

model in that it incorporates credit market frictions, which make external finance more expensive than internal finance. Furthermore, the premium on external finance depends on the collateral that a firm is able to provide, so that the cost of capital is directly related to asset prices. The more valuable are the assets that sit on a firm's balance sheet, the cheaper it is for the firm to access external finance.

A positive productivity shock has the usual positive effect on output and employment via improvements in the productivity of capital. However, the increase in asset prices associated with the shock also reduces the cost of capital, further increasing investment, output and asset prices. This positive feedback loop, known as the financial accelerator, means that the effects of a shock are amplified. It also gives extra traction to monetary policy. For example, a monetary loosening reduces real interest rates, raising asset prices, and therefore makes external finance less expensive.

The asset price bubble that Bernanke and Gertler add into this model is exogenously determined. They define the bubble as the difference between the market price of capital, S_t , and the fundamental value of capital, Q_t , which is simply the discounted sum of future dividends. The bubble develops according to the following rule:

$$\begin{aligned}
 S_{t+1} - Q_{t+1} &= \frac{a}{p} (S_t - Q_t) R_{t+1}^q && \text{with probability } p \\
 S_{t+1} - Q_{t+1} &= 0 && \text{with probability } (1 - p)
 \end{aligned}$$

where $p < a < 1$. In other words, the bubble grows until such time as it bursts, but the expected discounted value of the bubble decays over time. When the bubble bursts, the asset price goes instantaneously back to the fundamental value.

Bernanke and Gertler simulate their model under two alternative monetary policy rules. The first of these rules has the interest rate responding only to inflation, but under the second rule policy also reacts to the lagged asset price. They conclude that the best policy is to focus aggressively² on inflation and ignore asset prices, in that this policy achieves the lowest variance of output and inflation. Their simulations show that a monetary policy rule which accommodates inflation but responds to asset prices actually leads to a decline in output and inflation during a positive bubble. The rise in interest rates in response to the bubble drives down fundamental values to a greater extent than the bubble stimulates them. When the monetary policy rule aggressively targets inflation, they find that adding in the response to asset prices makes little difference, though what difference it does make is still destabilising. They conclude that:

a monetary regime that focuses on asset prices rather than fundamentals may well be actively destabilising. The problem is that the central bank is targeting the wrong indicator.

²They test two different parameters on expected inflation in the Taylor rule - one of which represents an aggressive response to inflation, whilst the other is more accommodating.

Cecchetti et al (2000 [27]) reach very different conclusions from the same model. They criticise Bernanke and Gertler for considering too narrow a set of Taylor rules and for failing to consider different parameterisations of the New Keynesian Phillips' curve. They report the results of simulations of the model in which they loosen these restrictions. In particular, they consider:

- (1) Taylor rules which include the output gap, and which allow for interest rate smoothing.
- (2) policy rules that react to asset mis-pricing rather than to the asset price itself. In other words, they assume that the central bank can distinguish between asset price movements caused by changes to the fundamentals, and those which are caused by a bubble.
- (3) the implications of making agents more or less backward looking in their wage setting. In other words, they vary the weights on past inflation and future expected inflation in the New Keynesian Phillips' curve.

They report that in the majority of cases, it is optimal for interest rates to respond to asset mis-pricing.

They further criticise Bernanke and Gertler on the basis that both the bubble's size and duration, and the level of leverage in the economy are treated as exogenous. They argue that when private agents expect the monetary authorities to 'prick' a bubble, the bubble is less likely to appear in the first place. Alternatively, if the bank can tighten policy in the formative stage of the bubble it will mitigate the

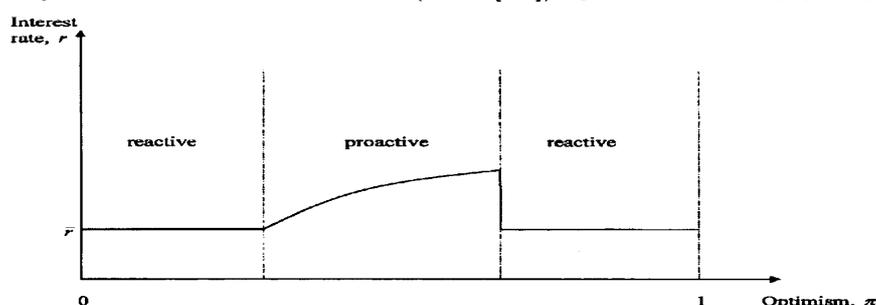
worst excesses that might otherwise occur. The authors also contend that if it is known that the monetary authorities will react to asset prices, then firms and households will react to stock market buoyancy by reducing their leverage, and this will dampen the effect of the financial accelerator.

Bordo and Jeanne (2002 [23]) suggest that the best way to think of asset price targeting is as costly insurance against financial crisis. In their highly stylised model, they incorporate a financial shock whose distribution depends on firms' indebtedness. The justification for the endogeneity of this shock is similar to Bernanke and Gertler's explanation of the financial accelerator. It lies in the fact that financial intermediaries rely on collateral to reduce financial frictions. Collateral in turn is driven by asset prices. Given that monetary policy can affect asset prices, and thereby debt accumulation, it also affects the probability of a damaging financial shock. A proactive monetary policy can thus prevent a credit crunch from emerging in the future. However, such a policy incurs a cost in terms of sacrificing short-run macroeconomic objectives.

The authors define indebtedness, D , as a decreasing function of monetary policy, r , and an increasing function of 'optimism', π :

$$D = D\left(\overset{+}{\pi}, \bar{r}\right)$$

Figure 4.1. Bordo and Jeanne (2002 [23]) optimal monetary policy



where optimism reflects the subjective probability associated with high future profits. They find that the optimal monetary policy depends decisively upon the optimism of the private sector. This result is illustrated by figure 4.1, which is lifted from their paper. When optimism is low, firms do not leverage themselves very highly, so the risk of a credit crunch is low, and the cost of a proactive policy is not worth bearing. As the private sector becomes more optimistic they increase their leverage and the probability of a credit crunch increases. It becomes worthwhile to insure against that risk with a proactive monetary policy. However, as optimism increases, there is also an increase in the cost of the proactive monetary policy. The cost increases because the more optimistic private agents are, the greater the interest rate that needs to be set to curb their indebtedness. At some point, the cost associated with distorting monetary policy becomes so high that it no longer pays to insure against the credit crunch.

In this way, Bordo and Jeanne conclude that there is no simple rule as to how central banks should respond to asset prices. The optimal policy depends on the

economic circumstances in a complex, non-linear way that cannot be represented in a straightforward Taylor rule.

Bean (2004 [15]) examines the effects of targeting asset prices within a simple New Keynesian model. His key conclusion is that expectations of future policy actions are at least as significant as current policy in preventing asset booms and busts. In his model, credit crunches occur with a given probability, but their severity depends upon the level of indebtedness in the economy. In the model, higher interest rates reduce capital formation and associated indebtedness, but the higher interest payments exactly offset this so that the output cost of a credit crunch is unaffected. In this way, current monetary policy does not have any impact on the severity of a credit crunch. However, monetary policy can effect the severity of future credit crunches through its impact on future expected output, and therefore on current capital accumulation and leverage. Hence, a central bank may find it optimal to use monetary policy commitments to limit the build up of leverage in the economy. The optimal commitment is in fact to stabilise output by less than the discretionary optimum when a credit crunch occurs. This counter-intuitive result arises because, by committing to a larger output cost if a credit crunch does occur, the central bank is disciplining private agents to limit their indebtedness.

In summary, a number of themes recur within the literature:

- (1) There is great difficulty in identifying whether asset price movements are driven by changes in the fundamentals or by noise trading. It is only with the benefit of hindsight that bubbles become recognisable. For many authors (see for example Greenspan 2002 [49]) this provides an overwhelming reason for not attempting to target asset prices. Cecchetti et al (2000 [27]), on the other hand, make an analogy between the concept of a fundamental asset price and the concept of potential output. They argue that measuring asset mis-pricing is of a similar complexity as measuring the output gap.
- (2) The macroeconomic consequences of bubbles are relatively mild in the absence of some kind of financial accelerator effect. As Bean (2004 [15]) states, “if the only macroeconomic consequences of booms and busts in asset prices were via conventional wealth effects on aggregate demand, then they would constitute little more than a nuisance to monetary policy makers”. It is only when falling asset prices combine with financial market frictions to cause credit rationing and credit crunches that significant welfare losses occur.
- (3) A number of authors argue that the magnitude of the monetary policy response that would be needed to correct for a bubble would risk causing serious harm to the real economy. Greenspan (2002 [49]) provides a selection of empirical evidence that suggests that the response of asset prices

to monetary policy is weak. Assenmacher-Wesche and Gerlach (2008 [9]) estimate VARs in order to assess the responses of equity and house prices to monetary policy across 17 different countries. They concur that using monetary policy to offset asset price movements in an attempt to guard against financial instability may have large effects on economic activity.

- (4) Conversely, Bean (2004 [15]) highlights the way in which a commitment to future policy may have significant effects on the expectations, and hence the behaviour, of the private sector. Such commitments, if they are effective in preventing bubbles from occurring, may never actually have to be acted upon. Cecchetti et al (2000 [27]) illustrate this in a simulation of the Bernanke and Gertler model. They compare Taylor rules with and without a response to asset prices. Although the asset targeting rule involves a larger response ex-ante to bubbles, ex-post the monetary policy response is smaller because private agents fully expect the central bank's response, and so bubbles do not grow as large.
- (5) Even if it is appropriate to target asset mis-pricing, the timing of monetary policy poses significant difficulties. The lags in the transmission of monetary policy mean that it may be counter-productive to respond to a bubble with a monetary tightening. If the bubble bursts of its own accord, just as the monetary tightening takes effect, then the economy will be hit simultaneously by two deflationary forces. Gruen et al (2003 [50]) show

that the informational requirements for implementing an asset price targeting policy are particularly stringent when these timing considerations are taken into account.

In the remainder of this chapter I will present a model which will attempt to address the final three themes highlighted above. Much has been said on the first theme, and though I will return to the issue of measuring mis-pricings in the conclusion, this is not an issue for theoretical modelling. As far as the second theme is concerned, the importance of financial accelerator effects is well established and relatively uncontroversial. Again, therefore, I will not concern myself with this issue in the model that follows, but will return to the issue in my conclusion.

My main concern in what follows is to provide a new perspective on the asset price targeting debate. All of the models discussed above either treat the bubble process as exogenous or use some simple but poorly specified construct to endogenise the bubble process. My main contribution is to use a specific behavioural framework to generate an endogenous bubble.

A further contribution is that I use a fully specified dynamic stochastic general equilibrium (DSGE) model to assess different policy rules. This allows me to do a full welfare analysis, rather than having to resort to ad hoc assessments using central bank loss functions.

4.2. The Model

I set up a standard New Keynesian DSGE model but inhabit it with investors who base their portfolio decision on one of two heuristical forecasting rules.

The model economy consists of a set of households, a set of firms and a bond-issuing government. The households derive an income by providing a differentiated labour service and consume a mixed bundle of output. They allocate their wealth intertemporally by holding government bonds or through capital ownership. Firms set the price for their output via Calvo contracts and households set the wage for their differentiated labour similarly. The nominal rigidity introduced by the Calvo mechanism is essential in providing a role for monetary policy.

I use the artifice of a perfectly competitive bundler to transform the differentiated output of the firms into a homogenous output-bundle which is consumed by the households or reinvested as capital. I use a similar bundler to transform the differentiated labour into a homogenous labour-bundle which is used in the productive process by the firms. Each of these firms pays a wage to labour and a rent to its capital owners.

To this extent I have a standard New Keynesian model (see, for example, Canzoneri, Cumby and Diba (2007[**26**])). My model only differs from this standard to the extent that agents' expectations of asset prices differ from the rational expectation. I model agents' beliefs about future asset prices in the same way as I did in the previous chapter. Each individual agent follows a behavioural process

whereby they choose between two simple heuristical forecasts, in the traditions of Frankel and Froot (1986 [42]) and De Grauwe and Grimaldi (2006 [31]).

4.2.1. Household Maximisation

There is a continuum of households indexed by h across the unit interval. Each household maximises its intertemporal utility function,

$$U_{h,t} = E_t \sum_{s=t}^{\infty} \beta^{s-t} \left\{ \ln(c_{h,s}) - \frac{l_{h,s}^{1+\phi}}{1+\phi} \right\}$$

where $c_{h,t}$ is the household's real consumption of the composite good, y_t , and $l_{h,t}$ is its differentiated labour supply. I consider a cashless economy in which households can transfer their wealth from one time period to another by holding government bonds, $b_{h,t}$, or investing in capital, $k_{h,t}$. Government bonds are one period, paying a pre-announced gross nominal return of $(1 + R_t^b)$. Capital is bought at the price P_t and has a rental return R_t^k each subsequent period. Capital depreciates at the rate δ and there is a capital adjustment cost. Hence, each household faces the budget constraint:

$$k_{h,t}R_t^k + b_{h,t}(1 + R_t^b) + l_{h,t}W_{h,t} = c_{h,t}P_t + i_{h,t}P_t + b_{h,t+1}$$

and the law of motion of capital is given by:

$$k_{h,t+1} = (1 - \delta) k_{h,t} + i_{h,t} - \frac{\nu}{2} \left(\frac{i_{h,t}}{k_{h,t}} - \delta \right)^2 k_{h,t}$$

where the final term represents the cost of capital adjustment.

This optimisation yields some standard results. Firstly, the marginal rate of substitution between consumption and leisure is set equal to a mark up on the real wage, because of imperfect competition in the labour market:

$$MRS_{c,t} = \frac{l_{h,t}^\phi}{c_{h,t}^{-1}} = \frac{W_{h,t}}{P_t} (1 + \varepsilon_{Lh})$$

where $\varepsilon_{Lh} = \frac{l_{h,t}}{W_{h,t}} \frac{\partial W_{h,t}}{\partial l_{h,t}}$ is the elasticity of demand for labour from household h .

Secondly, the expected real intertemporal returns on the different assets are equalised. When the return on bonds is discounted at the stochastic discount rate we get unity:

$$\frac{\beta E_t \Lambda_{h,t+1}}{\Lambda_{h,t}} (1 + R_{t+1}^b) = 1$$

Λ_t is the marginal utility of nominal wealth. The capital holding conditions are more complex because, in this case, the total return is a combination of the rental rate, the depreciation and the capital adjustment cost. However, I can still show that the expected discounted return from investing one dollar in capital is again

unity:

$$\frac{\beta E_t \Lambda_{h,t+1}}{\Lambda_{h,t}} \left\{ \frac{R_{t+1}^k}{P_t} + \frac{1 - \nu \left[\frac{i_{h,t}}{k_{h,t}} - \delta \right]}{1 - \nu \left[\frac{i_{h,t+1}}{k_{h,t+1}} - \delta \right]} \frac{P_{t+1}}{P_t} \left[(1 - \delta) + \frac{\nu}{2} \left[\left(\frac{i_{h,t+1}}{k_{h,t+1}} \right)^2 - \delta^2 \right] \right] \right\} = 1$$

It is also straightforward to show that the marginal rate of substitution between consumption this period and next is equal to the product of the intertemporal discount rate and the gross real interest rate:

$$\frac{E_t C_{h,t+1}}{C_{h,t}} = \beta (1 + E_t r_{t+1}^b)$$

where $1 + E_t r_{t+1}^b \equiv (1 + R_{t+1}^b) \frac{P_t}{E_t P_{t+1}}$.

4.2.1.1. Calvo wage setting. I introduce nominal rigidities in the labour market by using Calvo wage contracts. A proportion of households, $(1 - \omega)$, are free to adjust their wage. They choose the wage which maximises their utility across the states of nature for which that wage rate will hold. I assume that the remainder of the households simply update their last period wage by the steady state gross inflation rate, $(1 + \Pi)$.

Households sell their differentiated labour in a monopolistically competitive market to a perfectly competitive bundler. The bundler combines the labour of the various households into aggregate labour which is employed by the firms. The

bundling technology is a Dixit-Stiglitz aggregator:

$$l_t = \left[\int_0^1 l_{h,t}^{\frac{\gamma-1}{\gamma}} dh \right]^{\frac{\gamma}{\gamma-1}}$$

The bundler's cost minimisation implies that each household faces the following demand for their labour service:

$$l_{h,t} = \left(\frac{W_{h,t}}{W_t} \right)^{-\gamma} l_t$$

Households which are free to optimise in period t choose the wage rate, $W_{h,t}^*$, which maximises utility across the states of nature for which that wage rate will hold, subject to the labour demand curve, the budget constraint and the law of motion of capital. In other words, it maximises:

$$E_t \sum_{s=t}^{\infty} (\omega\beta)^{s-t} \left\{ \begin{array}{l} \ln(c_{h,s}) - \frac{[(1+\Pi)^{s-t} W_{h,t}^*]^{-\gamma} W_s^\gamma l_s^{1+\phi}}{1+\phi} \\ -\mu_{h,s} \left[k_{h,s+1} - (1-\delta)k_{h,s} - i_{h,s} + \frac{\nu}{2} \left(\frac{i_{h,t}}{k_{h,t}} - \delta \right)^2 k_{h,t} \right] \\ +\Lambda_{h,s} \left[\begin{array}{l} k_{h,s} R_s^k + ((1+\Pi)^{s-t} W_{h,t}^*)^{1-\gamma} W_s^\gamma l_s + \\ b_{h,s} (1+R_s^b) - c_{h,s} P_s - i_{h,s} P_s - b_{h,s+1} \end{array} \right] \end{array} \right\}$$

Therefore:

$$W_{h,t}^* = \left[\frac{\gamma}{\gamma - 1} \frac{E_t \sum_{s=t}^{\infty} (\omega\beta)^{s-t} (1 + \Pi)^{-\gamma(1+\phi)(s-t)} W_s^{\gamma(1+\phi)} l_s^{1+\phi}}{E_t \sum_{s=t}^{\infty} (\omega\beta)^{s-t} \Lambda_{h,s} W_s^{\gamma} l_s (1 + \Pi)^{(1-\gamma)(s-t)}} \right]^{\frac{1}{1+\gamma\phi}}$$

When wages are fully flexible (i.e. $\omega = 0$ and $W_t = W_{h,t}^*$) this reduces to $W_{h,t}^* = \frac{\gamma}{\gamma-1} \frac{l_t^\phi}{\Lambda_{h,t}}$. In other words, the wage is a mark up over the disutility of work.

The aggregate wage is given by:

$$\begin{aligned} W_t &= \left[\int_0^1 W_{h,t}^{1-\gamma} dh \right]^{\frac{1}{1-\gamma}} \\ &= \left[\sum_{i=0}^{\infty} (1 - \omega) \omega^i (\Pi^i W_{h,t-i}^*)^{1-\gamma} \right]^{\frac{1}{1-\gamma}} \\ &= \left[(1 - \omega) W_{h,t}^{*1-\gamma} + \omega \Pi^{1-\gamma} W_{t-1}^{1-\gamma} \right]^{\frac{1}{1-\gamma}} \end{aligned}$$

4.2.2. Firm Optimisation

There is a continuum of retail firms indexed by f across the unit interval. Each firm hires bundles of labour at the aggregate wage rate, W_t . They hire capital from the households in a perfectly competitive factor market at the rental rate of

capital, R_t^k . The firms make the decision of how much labour, $l_{f,t}$, and how much capital, $k_{f,t}$, to employ; and thus how much output to produce, $y_{f,t}$. They sell their output in a monopolistically competitive market at the price, $P_{f,t}$, and are constrained by production technology.

Each firm chooses an input mix to maximise profits:

$$y_{f,t}P_{f,t} - k_{f,t}R_t^k - l_{f,t}W_t$$

subject to its production function:

$$y_{f,t} = z_t k_{f,t}^\alpha l_{f,t}^{(1-\alpha)}$$

Note that the technology, z_t , is common across all firms.

The optimal inputs are therefore:

$$l_{f,t} = \frac{(1-\alpha)y_{f,t}P_{f,t}}{W_t}$$

$$k_{f,t} = \frac{\alpha y_{f,t} P_{f,t}}{R_t^k}$$

When firms charge different prices, the optimal level of inputs varies across firms.

However, the optimal capital to labour ratio is constant across firms:

$$\frac{k_{f,t}}{l_{f,t}} = \frac{\alpha}{(1-\alpha)} \frac{W_t}{R_t^k}$$

Because of this symmetry, marginal cost is also constant across firms. It can be derived as:

$$MC_{f,t} = \frac{W_t^{1-\alpha} R_t^{k\alpha}}{z_t \alpha^\alpha (1-\alpha)^{(1-\alpha)}}$$

I assume Calvo-type nominal rigidities in the goods market, similar to those in the labour market. In each period a randomly chosen fraction, $(1-\eta)$, of the firms are free to reset their prices. These firms set new prices taking their respective demand curves as given. The remainder of the retail firms cannot reoptimise, but adjust their price by the steady state inflation, $(1+\Pi)$.

If a firm has the opportunity to reset its price then it chooses the new price, $P_{f,t}^*$. The general price level is:

$$\begin{aligned} P_t &= \left[\int_0^1 P_{f,t}^{1-\theta} df \right]^{\frac{1}{1-\theta}} \\ &= \left[\sum_{i=0}^{\infty} (1-\eta)\eta^i (\Pi^i P_{f,t-i}^*)^{1-\theta} \right]^{\frac{1}{1-\theta}} \\ &= \left[(1-\eta)P_{f,t}^{*1-\theta} + \eta\Pi^{1-\theta}P_{t-1}^{1-\theta} \right]^{\frac{1}{1-\theta}} \end{aligned}$$

I assume the artifice of a perfectly competitive goods bundler employing Dixit-Stiglitz technology. Each individual firm, therefore, faces the demand curve:

$$y_{f,t} = \left(\frac{P_{f,t}}{P_t} \right)^{-\theta} y_t$$

In periods when the firm gets the opportunity to choose a new price, it chooses the price which maximises its expected discounted future stream of profits across the states of nature for which that price will hold. In other words, it maximises:

$$E_t \sum_{s=t}^{\infty} (\eta\beta)^{s-t} \frac{\Lambda_s}{\Lambda_t} \left[y_s P_s^\theta \{ (1 + \Pi)^{s-t} P_{f,t}^* \}^{1-\theta} - TC_{f,s} \right]$$

This yields the optimal price:

$$P_{f,t}^* = \frac{\theta}{\theta - 1} \frac{E_t \sum_{s=t}^{\infty} (\eta\beta (1 + \Pi)^{-\theta})^{s-t} \Lambda_s P_s^\theta y_s MC_{f,s}}{E_t \sum_{s=t}^{\infty} (\eta\beta (1 + \Pi)^{1-\theta})^{s-t} \Lambda_s P_s^\theta y_s}$$

In other words, the firm sets the price so that its expected value is equal to a mark up, $\frac{\theta}{\theta-1}$, over expected marginal cost. In the case of no price stickiness (i.e. $\eta = 0$), $P_{f,t}^* = \frac{\theta}{\theta-1} MC_{f,t}$. This is the standard result that under monopolistic competition firms set price as a mark up over marginal cost.

4.2.3. Aggregation

For tractability, I assume that there are full contingent claims markets. Given the ex-ante homogeneity of the households, this ensures that consumption and wealth are constant across all households. Effectively, risk averse households will insure against not being able to adjust their wage rate. Hence, $c_{h,t} = c_t \forall h$ and $\Lambda_{h,t} = \Lambda_t \forall h$.

It also entails that all households which are free to optimally set their wage in a given period are in exactly the same position and will choose the same wage, $W_{h,t}^* = W_t^* \forall h$.

Firms which are free to set their optimal price are also all in identical positions, and so $P_{f,t}^* = P_t^* \forall h$. Furthermore, I have already shown that marginal cost and the capital to labour ratio are the same across all firms.

In order to aggregate output, I begin by noting that an individual firm's demand and supply must be equal, and then integrating across all firms:

$$y_{f,t} = z_t k_{f,t}^\alpha l_{f,t}^{(1-\alpha)} = P_t^\theta P_{f,t}^{-\theta} y_t$$

$$\int_0^1 z_t k_{f,t}^\alpha l_{f,t}^{(1-\alpha)} df = \int_0^1 P_t^\theta P_{f,t}^{-\theta} y_t df$$

$$\int_0^1 z_t \left(\frac{k_{f,t}}{l_{f,t}} \right)^\alpha l_{f,t} df = \int_0^1 P_t^\theta P_{f,t}^{-\theta} y_t df$$

$$z_t \left(\frac{k_t}{l_t} \right)^\alpha \int_0^1 l_{f,t} df = y_t P_t^\theta \int_0^1 P_{f,t}^{-\theta} df$$

$$y_t = \frac{z_t k_t^\alpha l_t^{(1-\alpha)}}{pd_t}$$

where $pd_t = P_t^\theta \int_0^1 P_{f,t}^{-\theta} df$

In other words, aggregate output is a decreasing function of price dispersion, pd_t .

4.2.4. Asset Prices

The model so far expounded is a relatively standard New-Keynesian model. In this section I depart from that standard and introduce heterogeneous forecasting of asset prices. The behavioural finance literature suggests that the assumption of rational expectations in asset markets is difficult to support. The suggestion is that asset markets are prone to uncertainty and speculation in a way in which goods markets and labour markets are not. For this reason a more complex specification of the forecasting rules employed in asset markets is needed, rather than a simple appeal to rationality.

I will, however, begin with an account of asset prices under rational expectations. This will provide a benchmark against which to assess the behavioural

model which I consider thereafter. It also provides us with a measure of the fundamental value of the asset, and this will form the basis for one of the behavioural forecasting rules.

4.2.4.1. Asset Prices Under Rational Expectations. If I were to model an equity market in a rational expectations framework, I could simply introduce equity trading into the households' optimisation problem. This would entail an equity price which is the discounted sum of all future rental payments to capital:

$$Q_t^* = E_t \sum_{s=t+1}^{\infty} \beta^{s-t} \frac{\Lambda_s}{\Lambda_t} R_s^k$$

Alternatively, I can express this as the discounted value of the sum of the next period rental payment and price:

$$Q_t^* = E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t} [R_{t+1}^k + Q_{t+1}^*]$$

I define Q_t^* , the price of equity in the rational model, as the fundamental value of equity.

Equivalently, in real terms, I have:

$$(4.1) \quad q_t^* = E_t \beta \frac{\lambda_{t+1}}{\lambda_t} [r_{t+1}^k + q_{t+1}^*]$$

In this way, I can see that in a rational expectations model the present equity price is a function of future expectations of the equity price, or equivalently of the

infinite series of returns. It is the ability of agents to form such expectations in highly volatile equity markets that behavioural finance questions.

Behavioural finance suggests that, because of the limited rationality of individual agents and effective limits on arbitrage, the actual asset price at any time can deviate from the fundamental level that q_t^* represents. This is equivalent in a macroeconomic context to arguing that there is a bias to the expected future return on capital. In the next section, we derive that bias.

4.2.4.2. Asset Prices Under Behavioural Assumptions. I follow in the tradition of Frankel and Froot (1986 [42]) and De Grauwe and Grimaldi (2006 [31]) in specifying an alternative determinant of asset prices. They highlight the importance of the frictions between chartist and fundamental forecasts in the determination of asset prices.

I assume that agents operating in the capital markets choose between two simple heuristical rules when forecasting future asset prices. At the beginning of each period chartist and fundamentalist forecasts of the asset price this period and next are formed. The forecasts of this period's asset price then determine the actual asset price via a bargaining process. The forecasts of next period's asset price imply a particular expected return on capital.

The chartist forecasts of the present and next period asset price are a simple extrapolation of the historical price series:

$$(4.2) \quad E_{c,t}(q_t) = q_{t-1} + \chi_c(q_{t-1} - q_{t-2}) + \chi_c^2(q_{t-2} - q_{t-3}) + \dots$$

$$(4.3) \quad E_{c,t}(q_{t+1}) = E_{c,t}(q_t) + \chi_c(E_{c,t}(q_t) - q_{t-1}) + \chi_c^2(q_{t-1} - q_{t-2}) + \dots$$

It is arguable as to whether such heuristical rules should be specified in real or nominal terms. I choose real terms on the basis that, in this paper, I am attempting to address the issue of asset market bubbles, and hence I want to avoid the issue of money illusion.

The fundamentalist forecast is that the asset price will move back towards its fundamental value, q_t^* , during the next period, unless it is already close to the fundamental, defined by the bounds $\pm C$:

$$(4.4) \quad \begin{aligned} E_{f,t}(q_t) &= q_{t-1} - \chi_f(q_{t-1} - q_t^*) && \text{where } |q_{t-1} - q_t^*| > C \\ &= q_{t-1} && \text{where } |q_{t-1} - q_t^*| \leq C \end{aligned}$$

$$(4.5) \quad \begin{aligned} E_{f,t}(q_{t+1}) &= E_{f,t}(q_t) - \chi_f(E_{f,t}(q_t) - q_t^*) && \text{where } |q_{t-1} - q_t^*| > C \\ &= E_{f,t}(q_t) && \text{where } |q_{t-1} - q_t^*| \leq C \end{aligned}$$

The actual asset price is determined via a bargaining process between those who favour the chartist rule and those who favour the fundamentalist forecast:

$$(4.6) \quad q_t = w_{c,t} E_{c,t}(q_t) + w_{f,t} E_{f,t}(q_t)$$

where the weights depend on the past performance of the two forecasts. I use the functional form:

$$(4.7) \quad w_{c,t} = \frac{\exp(v\Omega_{c,t})}{\exp(v\Omega_{f,t}) + \exp(v\Omega_{c,t})}$$

$$(4.8) \quad w_{f,t} = \frac{\exp(v\Omega_{f,t})}{\exp(v\Omega_{f,t}) + \exp(v\Omega_{c,t})}$$

The parameter v represents the propensity with which agents switch between forecasting rules. $\Omega_{c,t}$ and $\Omega_{f,t}$ are the excess returns over holding bonds associated with following the chartist forecast and the fundamentalist forecast respectively. They are calculated as follows:

$$(4.9) \quad \Omega_{c,t} = [q_{t-1} - q_{t-2} (1 + r_{t-1}^b)] \cdot \text{sign} [E_{c,t-2}(q_{t-2}) - q_{t-2}]$$

$$(4.10) \quad \Omega_{f,t} = [q_{t-1} - q_{t-2} (1 + r_{t-1}^b)] \cdot \text{sign} [E_{f,t-2}(q_{t-2}) - q_{t-2}]$$

where $q_{t-2} (1 + r_{t-1}^b)$ represents the return to investing funds in bonds and q_{t-1} is the return to investing the same funds in equities. $\text{sign} [E_{c,t-2}(q_{t-2}) - q_{t-2}]$ takes the value -1 when $E_{c,t-2}(q_{t-2}) < q_{t-2}$, in which circumstances an agent following the chartist rule would choose to invest in bonds rather than equities, and takes

the value +1 when chartists choose equities over bonds. $sign [E_{f,t-2} (q_{t-2}) - q_{t-2}]$ is analogous.

The average expected return on capital that is implied by these behavioural forecasting rules can be solved for by substituting the actual asset price and the weighted average forecast of next period's asset price for the respective fundamental values in equation 4.1. Therefore:

$$(4.11) \quad E_{b,t} r_{t+1}^k = \frac{q_t}{\beta} \frac{\lambda_t}{E_t \lambda_{t+1}} - [w_{c,t} E_{c,t} (q_{t+1}) + w_{f,t} E_{f,t} (q_{t+1})]$$

In this way, we can think of behavioural rules as driving an asset mis-pricing by inducing a bias in the expected future return on capital.

4.2.5. The Government

The government's only role in this model is as a bond issuer. I assume that the government sets the nominal interest rate, R_{t+1}^b , according to a Taylor rule which includes a response to the most recent asset mis-pricing:

$$(4.12) \quad R_{t+1}^b = r^{b*} + \Pi_t + \zeta_{\Pi} (\Pi_t - \Pi^*) + \zeta_Y \ln \left(\frac{y_t}{y^*} \right) + \zeta_Q \ln \left(\frac{q_{t-1}}{q_{t-1}^*} \right)$$

where r^{b*} and y^* are the steady state real interest rate and output respectively; Π^* is the target inflation rate; and q_{t-1}^* is the fundamental asset price. The government supplies as many bonds as are demanded at this interest rate.

The response is to the asset mis-pricing in period $t - 1$ because we are in a world in which asset prices cannot easily be predicted. In rational models, policy can react to current variables, which in turn depend upon policy, because all variables are determined simultaneously. There is an implicit assumption that agents costlessly form entire response functions and costlessly and instantaneously adjust their trading volumes in response to price signals. This is inconsistent with the essence of behavioural economics. In the behavioural world, the government cannot perfectly anticipate how private agents will respond to its policy prescriptions. Hence, the simultaneous realisation of monetary policy and asset prices is not within the spirit of a behavioural model. Equivalently, a solution method for such a system of equations would require the imposition of some concept of rational consistency. The Taylor rule that is most compatible with the spirit of behavioural modelling, therefore, is one in which the monetary authorities react to the mis-pricing from the previous period.

The aim in this paper is to assess whether the central bank should take account of asset prices in setting monetary policy. I will do this by comparing the welfare effects of various parameterisations of the weight on the asset price, ζ_Q .

4.2.6. Welfare

I use a strictly utilitarian notion of welfare, defining it as aggregate utility:

$$\begin{aligned}
U_t &= \int_0^1 U_{h,t} \, dh \\
&= E_t \sum_{s=t}^{\infty} \beta^{s-t} \left\{ \int_0^1 \ln(c_{h,s}) \, dh - \frac{\int_0^1 l_{h,s}^{1+\phi} \, dh}{1+\phi} \right\}
\end{aligned}$$

Given our assumption of complete contingent claims markets, consumption is constant across all households. Therefore, the aggregate (or average) utility derived from consumption is just the same as the utility of consumption for any individual household. However, given price stickiness, firms employ different amounts of labour from different households, and so the aggregate disutility of labour is not straightforwardly related to the disutility of an individual household. I calculate it as follows:

$$\begin{aligned}
\frac{\int_0^1 l_{h,s}^{1+\phi} \, dh}{1+\phi} &= \frac{\int_0^1 (W_t^\gamma W_{h,t}^{-\gamma} l_t)^{1+\phi} \, dh}{1+\phi} \\
&= \frac{l_t^{1+\phi}}{1+\phi} w d_t
\end{aligned}$$

where $w d_t = W_t^{\gamma(1+\phi)} \int_0^1 W_{h,t}^{-\gamma(1+\phi)} \, dh$ is a measure of wage dispersion.

Therefore, welfare is given by:

$$U_t = E_t \sum_{s=t}^{\infty} \beta^{s-t} \left\{ \ln(c_s) - \frac{l_s^{1+\phi} w d_s}{1+\phi} \right\}$$

We can clearly see that nominal rigidities have an adverse effect on welfare. Wage dispersion directly increases the aggregate disutility of work. Price dispersion, on

the other hand, indirectly reduces welfare by reducing aggregate output, and hence consumption.

4.2.7. Parameterising the Model

My intention in this chapter is to pose questions about the effectiveness of monetary policy as a tool for alleviating the damaging effects of asset price bubbles. It is not to produce a model for calibrating the optimal policy. I, therefore, make no attempt to estimate the model. Instead, I borrow my parameterisation of the model from previous work. The parameters for the New Keynesian aspects of the model are taken from Canzoneri et al (2007 [26]), whilst those for the behavioural aspects are taken from the estimation of the model in the previous chapter. The baseline parameterisation that I use is given in appendix B.1.

4.3. Solving the Model

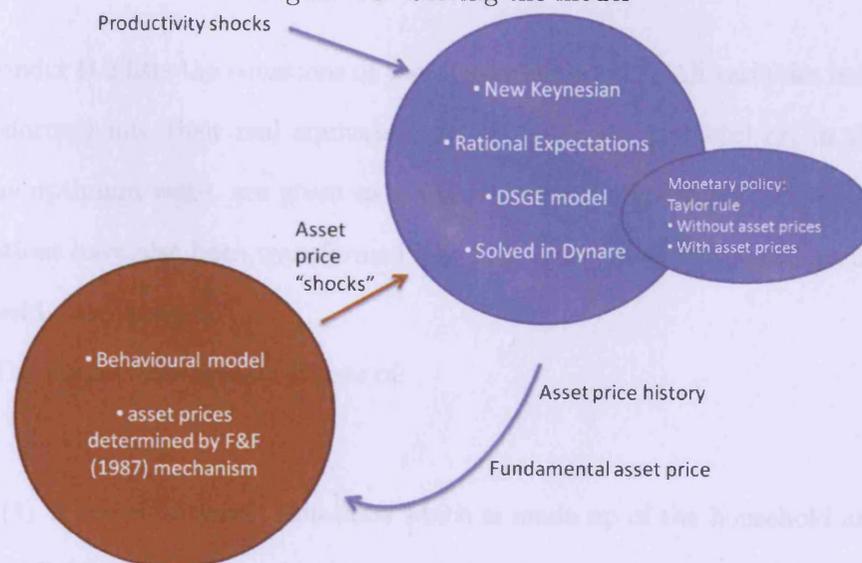
Solving this model presents very significant challenges. I am assuming rationality in the goods and labour markets, which means that agents are forward looking and understand the model, so that all of the equations that describe these markets must hold in both the present period and in their expected terms for all future periods. On the other hand, I am assuming that rationality breaks down in the asset market. In keeping with the spirit of the behavioural finance literature, agents make use of simple heuristical rules to determine what they consider to be a fair price for the asset and also in determining how the asset price will behave in the future. However, due to the complexity of the asset market, agents do not understand the behaviour of others in the market, and so the market clearing condition only holds in the present period. It does not hold in its expected future forms.

I can always impose some kind of asset mis-pricing. However, as long as agents understand the mechanism which causes that mis-pricing, and they rationally expect it to persist, then all that happens is that there is a persistent dislocation between the fundamental value of the asset and its price. If future dislocations are fully anticipated then the net effect is that the expected return to asset-holding is unchanged. If this is the case, then there will be no distortions in the wider economy. The key here is that it is biases to the expected return on asset holding, and not the mis-pricing of the asset in and of itself, which drives distortions in the allocation of resources.

My approach to solving the model is to take advantage of well established techniques for solving rational expectations models, and then introduce the non-rational aspects of the model via ‘shock’ processes. Shocks to the current and future expected asset price drive these variables away from their fundamental values. Of course, these are not shocks in the conventional sense. The behavioural model describes exactly how these deviations from fundamental are determined. As already stated, if the agents in the model understood this behavioural process then they could form a rational expectation of future deviations, and the asset mis-pricing would not effect the wider economy. By introducing these biases to the current and future asset price as unanticipated shocks, they cause a bias to the expected return on capital, and thereby alter the allocations throughout the economy.

Figure 4.2 summarises the solution method. I solve the majority of the model as if it were a rational expectations model. I do this using the Dynare pre-processor and the MATLAB programming environment. Dynare uses perturbation techniques to solve a system of non-linear equations with forward looking variables. The solution consists of a set of policy and transition functions which describe how each variable is determined by pre-determined variables (initial conditions) and a set of ‘shocks’. There are four ‘shocks’ to the ‘rational’ model. The first is a conventional productivity shock. There are a pair of asset price ‘shocks’. One affects the present asset price, driving a wedge between the actual asset price and its fundamental value. The other asset price ‘shock’ drives a wedge between the expected

Figure 4.2. Solving the model



future asset price and its fundamental value. Combined, these two shocks have the effect of biasing the expected return on capital, which in turn drives distortions in the rest of the model. The fourth 'shock' is the monetary policy response to the asset mis-pricing.

The productivity shock is treated as a random variable. The other three shocks, however, are endogenously determined. The asset price 'shocks' are determined by the behavioural asset pricing rules explained above (equations 4.2 to 4.10). The monetary policy response is determined by the Taylor rule in equation 4.12.

The various constituent parts of the model are described in greater detail below.

4.3.1. The Dynare Sub-Model

Appendix B.2 lists the equations of the dynare sub-model. All variables have been transformed into their real equivalent ($p_t^* = \frac{P_t^*}{P_t}$, $\lambda_t = \Lambda_t P_t$ etc) or, in the case of the optimum wage, are given as a ratio to the average wage ($w_t^* = \frac{W_t^*}{W_t}$). All equations have also been transformed into first order difference equations in order to avoid infinite sums.

The dynare sub-model consists of:

- (1) A set of ‘rational’ equations which is made up of the household and firm optimisations, production constraint, market clearing conditions, monetary policy rule, welfare definition and the rational asset pricing equation 4.1. The variables in this part of the model are denoted with an R superscript in the appendix.
- (2) A set of ‘behavioural’ equations that is made up of exactly the same features except:
 - the asset price, q_t^B , is set equal to its fundamental value (the asset price from the ‘rational’ equations, q_t^{*R}) plus an ‘asset price shock’ term, ε_t^q .
 - the expected future asset price, q_t^{BF} , is set equal to its fundamental value (the expected future asset price from the ‘rational’ equations, $E_t q_{t+1}^{*R}$) plus an ‘expected asset price shock’ term, ε_t^{qF} .

- the return on capital is derived from the asset price and expected future asset price via equation 4.13.

The variables in this part of the model are denoted with a B superscript in the appendix. Note that the expected future asset price, q_t^{BF} , is not a rational expectation of the value that the asset price, q_t^B , will take in the next period. This is where the lack of rationality enters the model, and it permeates the ‘behavioural’ equations by biasing the expected return on capital:

$$(4.13) \quad E_t r_{t+1}^{kB} = \frac{q_t^B}{\beta} \frac{\lambda_t^R}{E_t \lambda_{t+1}^B} - q_t^{BF}$$

- (3) A standard productivity process.

In this way, I have two, almost distinct, sets of equations. The ‘rational’ set, along with the productivity process, constitute an independent model in which all expectations are formed rationally. These equations do not depend upon the ‘behavioural’ equations in any way, and can be solved separately. The ‘behavioural’ equations, on the other hand require the fundamental asset price and rationally expected future asset price as pre-determined inputs if they are to be solved.

In fact, I need to include the ‘rational’ equations in the dynare model for two reasons. Firstly, the behavioural biases that come from the behavioural sub-model are calculated as deviations from the fundamental, and so I need to know what the fundamental is when I import the deviations as shocks into the dynare model.

Secondly, the fundamental value of the asset is the basis of the fundamentalist forecast in the behavioural sub-model, so I need to be able to calculate this and export it to the behavioural model.

The dynare sub-model can be solved to provide a set of functions which determine how that model's variables depend on predetermined variables and the four dynare 'shocks': the asset price 'shock', ε_t^q ; the expected asset price 'shock', ε_t^{qF} ; the policy response, ε_t^{MP} ; and the productivity shock, ε_t . In this way we have functions which we can iterate to find the time paths of variables in our model, once we have determined the size of the 'shocks'. These are determined by the behavioural sub-model, the Taylor rule and the productivity process.

The dynare code is shown in appendix B.3.

4.3.2. The Behavioural Sub-Model

The behavioural sub-model is a set of equations which, given the relevant initial conditions, solves to give the size of both the current and future expected asset mis-pricings. It consists of equations 4.2 to 4.10, transformed where appropriate into first order difference equations in order to avoid infinite sums. The behavioural sub-model also includes an equation which describes the monetary policy response to an asset mis-pricing.

The full listing of the equations is given in appendix B.4.

4.3.3. The Complete Model

In the next section, I will analyse some of the characteristics of the model. My ambition in this section is limited to describing the mechanics of the model and, in particular, how the dynare and behavioural sub-models mesh together. A simple way of illustrating this is to consider how a productivity shock is propagated through the model.

Solving for the steady state is straightforward. In the steady state, the asset price is equal to its fundamental value and is constant over time, so that both the fundamentalist and chartist forecasts are fulfilled. Hence, there is no effective behavioural bias. The steady state of the dynare sub-model can be solved in the usual way, by dropping time subscripts and solving the resulting set of simultaneous equations.

Beginning from steady state, a productivity shock impacts on the ‘rational’ variables in the dynare sub-model in the same way as we would expect in a standard New Keynesian model. In a fully rational model, one of the variables that would be affected is the asset price. A positive productivity shock increases the productivity of capital, and thereby also increases the returns to capital and the asset price. In this model, the fundamental price of the asset is impacted directly

by the productivity shock. The increase in the fundamental value drives up the fundamentalist forecast, and therefore the actual asset price as well.

In terms of the mechanics of the model, I first solve the dynare sub-model. This leaves me with a set of policy and transition functions which describe how each variable is determined by pre-determined variables (initial conditions) and the ‘shocks’. To simulate the model for a given realisation of the productivity shock, I proceed through the following routine:

- (1) For $t = 0$, set all variables equal to their steady state value.
- (2) Use the steady state values, and the realisation of the productivity shock in $t = 1$ to derive the fundamental value of the asset in the dynare sub-model. At this stage, assume the other three ‘shocks’ are equal to zero.
- (3) Use the behavioural sub-model to determine the size of the current and future expected asset mis-pricings, ε_t^q and ε_t^{qF} .
- (4) Re-solve the dynare sub-model for period 1, with the original productivity shock and the two asset mis-pricing ‘shocks’.
- (5) Solve the dynare sub-model for period 2 with the realisation of the productivity shock in that period; using the realisations of all the variables from period 1 as initial conditions, and setting $\varepsilon_t^q = \varepsilon_t^{qF} = 0$. Hence, derive the fundamental value of the asset in period 2.
- (6) Use this fundamental value, and the past realisations of the asset price to solve the behavioural sub-model. This time the behavioural sub-model

determines the policy response to the previous period's mis-pricing, as well as the two asset mis-pricing 'shocks'.

- (7) Re-solve the dynare sub-model for period 2, with the original productivity shock, the asset mis-pricing 'shocks', and the policy response.
- (8) Iterate steps 5 to 7 for each subsequent time period.

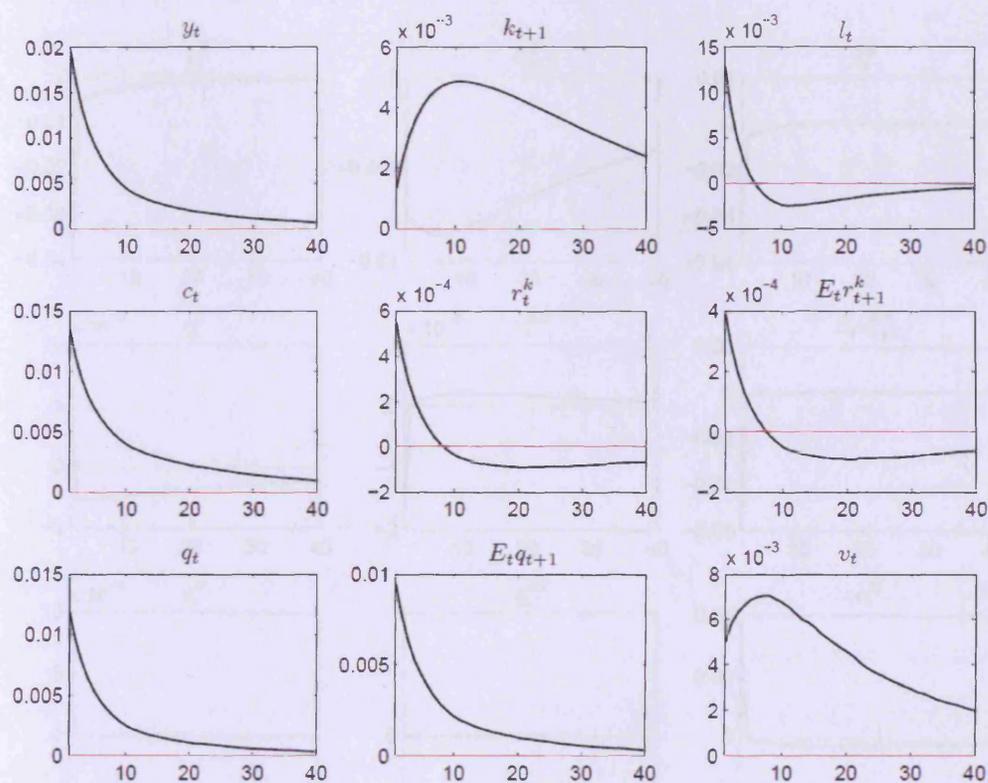
In this way, I can solve the full time-series of each variable. The assumption of rational expectations in the goods and labour markets is preserved, but decisions in the asset market are governed by behavioural rules. Similarly, the monetary policy response to inflation and the output gap are fully anticipated, but any response to an asset mis-pricing is unanticipated. Hence, the model is consistent with the general theme of the behavioural finance literature, that it is the particular complexity of financial markets that drives agents to use heuristical rules.

4.4. Preliminary Analysis of the Dynare sub-model

If I simply treat the asset price shocks and the policy response as random shocks, with mean zero, then the dynare sub-model collapses to the rational expectations paradigm. Furthermore, if I restrict them to always equal zero, then I have a rational model with a single technology shock. This provides a relevant benchmark against which to compare the results of our full model.

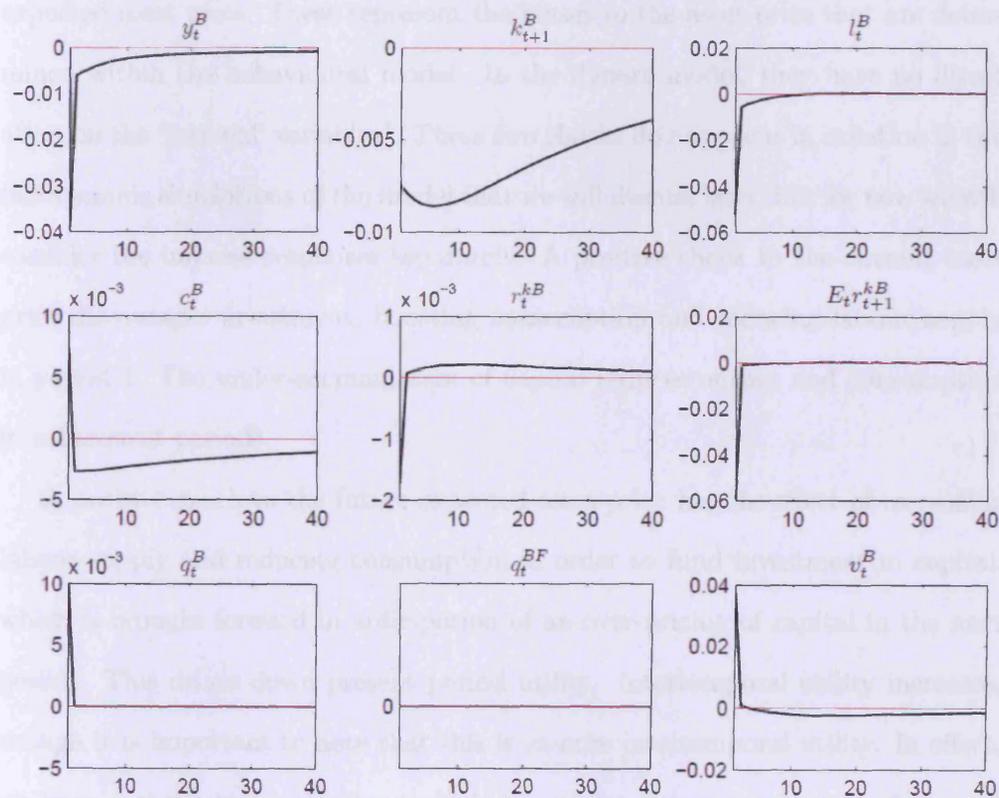
Here, I present and briefly analyse the impulse responses to all four ‘shocks’ in the dynare sub-model.

Figure 4.3 shows the responses of some key variables to a 1% productivity shock, with an auto-regressive component, ρ , of 0.95. Given that there are no asset price ‘shocks’, the behavioural variables in dynare follow exactly the same path as their ‘rational’ counterparts. The direct effect of the increase in productivity increases output, y , by 1% in period 1. Output actually rises by more than this because the productivity shock also affects the employment of capital and labour. It increases the productivity of both capital, k , and labour, l , which in turn lead to increases in investment and the demand for labour. The increase in permanent income drives up consumption, c , and reduces the supply of labour. In the periods immediately following the shock, the expansion in demand for labour outweighs the contraction in supply, but this is eventually reversed with the quantity of labour falling below its steady state value from around period 5. Discreet period utility, v , is driven

Figure 4.3. Impulse responses to a 1% productivity shock (i.e. $\varepsilon_1 = 0.01$)

mainly by changes in consumption, which are of an order of magnitude greater than the changes in labour. Initially, the return on capital and expected future return on capital are driven up by the productivity shock as the employment of labour increases rapidly, and the accumulation of capital lags behind. By period 9, however, sufficient capital has been accumulated so that the returns to capital have fallen back below their steady state level, as productivity falls back towards

Figure 4.4. Impulse responses to a 1% shock to the current asset price (i.e. $\varepsilon_1^q = 0.01$)



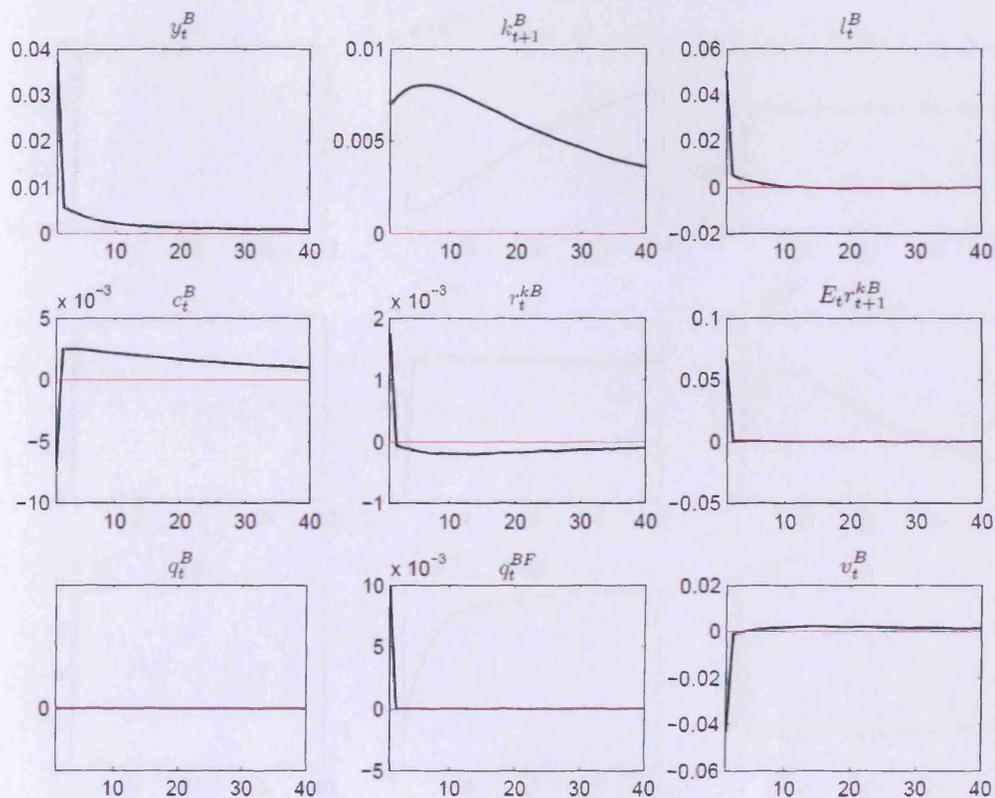
its steady state level. The asset price, q , shoots up in response to the productivity shock and gradually falls back to its steady state level. This is perfectly anticipated (see $E_t q_{t+1}$) since there are no other shocks after period 1.

Figure 4.4 shows the response to a transitory 1% shock to the current asset price, whilst figure 4.5 shows the response to a transitory 1% shock to the future expected asset price. These represent the biases to the asset price that are determined within the behavioural model. In the dynare model, they have no direct effect on the ‘rational’ variables³. These two shocks do not occur in isolation in the full dynamic simulations of the model that we will discuss later, but for now we will consider the impulse responses separately. A positive shock to the current asset price discourages investment, boosting consumption and reducing labour supply in period 1. The under-accumulation of capital reduces output and consumption in subsequent periods.

A positive shock to the future expected asset price has the effect of increasing labour supply and reducing consumption in order to fund investment in capital, which is brought forward in anticipation of an over-pricing of capital in the next period. This drives down present period utility. Intertemporal utility increases, though it is important to note that this is ex-ante intertemporal utility. In effect, we have a shock to expectations which is a wedge between ex-ante and ex-post returns to capital and which also drives a wedge between ex-ante and ex-post intertemporal utility. In fact, by the time agents reach period 2 and recognise their mistaken beliefs in period 1, they have already accumulated extra capital

³However, it is worth noting that they do affect the accumulation of capital in the full model, and hence affect the initial conditions for subsequent time periods. In this way, these shocks do have an effect on future fundamentals.

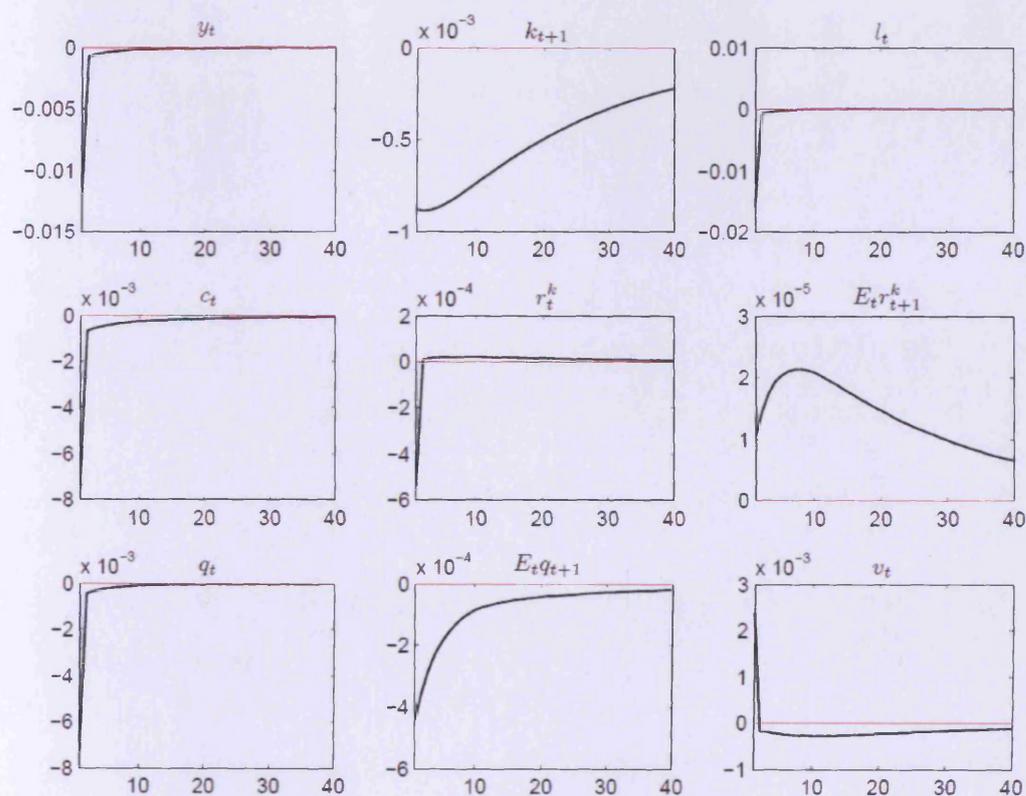
Figure 4.5. Impulse responses to a 1% shock to the expected future asset price (i.e. $\varepsilon_1^{q^F} = 0.01$)



which advantages them in period 2 and henceforth, though not to the extent that they had anticipated in period 1.

Figure 4.6 shows the response to a 1% transitory shock to the Taylor rule. As with the productivity shock, the monetary policy 'shock' in isolation does not

Figure 4.6. Impulse responses to a 1% shock to the Taylor rule (i.e. $\varepsilon_1^{MP} = 0.01$)



cause the rational and behavioural variables in the dynare model to diverge - it affects the equivalent variables in exactly the same way. If the nominal interest rate is set above its steady state value then, in the presence of price rigidities, the real return to bond-holding increases and agents substitute out of capital and into

bonds. The decrease in capital reduces the marginal productivity of labour, and hence labour falls. A fall in output, consumption and utility ensues.

4.5. Results

In order to assess the effect of including an asset price target in the monetary policy rules, I simulate the entire model under different parameterisations of the Taylor rule. The parameterisation that I use in the benchmark model is as follows:

$$R_{t+1}^b = 0.01 + \Pi_t + 2.02(\Pi_t - \Pi^*) + 0.184 \ln\left(\frac{y_t}{y^*}\right) + \zeta_Q \ln\left(\frac{q_{t-1}}{q_{t-1}^*}\right)$$

This is based on Canzoneri et al's (2007 [26]) estimation of a Taylor rule over the Volcker and Greenspan years as Federal Reserve chairmen (1979 - 2003). Of course, the estimated Taylor rule does not include a response to asset mis-pricing, so $\zeta_Q = 0$.

I run 1,000 stochastic simulations of the model under each of several different values of ζ_Q . I consider:

- (1) a 'passive' monetary regime, where there is no response to asset mis-pricings. In other words, $\zeta_Q = 0$.
- (2) a variety of 'activist' regimes with different weights, both positive and negative, on the asset mis-pricing. The weights I consider are $\zeta_Q = 0.05, -0.05, 0.1, -0.1, 0.5, -0.5, 1$ and -1 .

For each individual simulation, I run the model for forty periods. The only exogenous shock is the productivity shock, and this is drawn at random from a

normal distribution with mean zero and standard deviation 0.0086⁴ for each of the forty periods. I calculate the actual welfare each period and, for the final period, I also calculate the expected future welfare. I then discount these values to get a measure of inter-temporal welfare for each simulation. Averaging this welfare measure across the 1,000 simulations gives me a measure of expected welfare under each alternative monetary regime.

As is the case for any measure of welfare, the cardinal units are more or less meaningless. I follow the convention, initiated by Lucas (2003 [65]), of calculating and reporting consumption equivalents. The welfare measures for all of the results reported in this chapter refer to the proportion of consumption that households would be prepared to give up permanently, holding work effort constant, in order to live in a rational world with no asset bubbles and no associated monetary response.

Table 4.1 presents the estimation of welfare in the benchmark behavioural model, under different specifications of the Taylor rule. The second column of the table states the loss in welfare relative to the rational model under each monetary regime.

The behavioural model with a passive monetary regime ($\zeta_Q = 0$) leads to expected welfare which is equivalent to a loss of 0.177 of one percent of permanent consumption relative to the rational model. We can think of this as the cost of

⁴This parameter comes from an estimate of the 1960 - 2002 US data (with a log linear trend) by Canzoneri et al (2007 [26])

Table 4.1. Utility cost of various Taylor rules in the benchmark model

Monetary regime $\zeta_Q =$	Utility cost versus rational model (%age permanent consumption)	Performance relative to $\zeta_Q = 0$ (%age bias corrected)
0.00	-0.177 %	0.0 %
0.05	-0.187 %	-6.1 %
-0.05	-0.167 %	5.5 %
0.10	-0.199 %	-12.8 %
-0.10	-0.158 %	10.5 %
0.50	-0.357 %	-101.9 %
-0.50	-0.113 %	35.9 %
1.00	-1.083 %	-512.8 %
-1.00	-0.113 %	36.3 %

irrationality in the asset markets. I use this value as a basis against which to compare the performance of the alternative monetary policy options.

For example, the behavioural model with a modest ‘leaning against the wind’ monetary regime of $\zeta_Q = 0.05^5$ experiences a loss in welfare relative to the rational world which is equivalent to 0.187 of one percent of permanent consumption. Compared to the passive regime, this represents an exacerbation of the misallocation of resources which irrationality has caused. The welfare loss is 6.1% higher than under the passive monetary regime. This measure of the performance of activist monetary policy rules relative to the passive regime is reported in the third column of table 4.1.

⁵This is the monetary policy response recommended by Cecchetti et al (2000 [27])

In other words, a negative number in column 3 signals a monetary policy rule that increases the cost of behavioural biases - we can think of this as a monetary policy rule that exacerbates those biases. On the other hand, a positive number in column 3 signals a monetary policy rule that reduces the cost of behavioural biases - we can think of this as a monetary policy rule that corrects for those biases.

The surprising result is that 'leaning against the wind' policies (policies which lead to a monetary tightening when asset prices are above fundamental, and a loosening when they are below fundamental, i.e. $\zeta_Q > 0$) exacerbate the cost of behavioural biases. Even more surprisingly, the contrary 'running with the wind' policy, where an asset bubble is met with a monetary loosening, actually ameliorates the effects of the bubble. An explanation of these counter-intuitive results is required, and I will provide this in the next section.

Before providing that explanation, I shall briefly consider the significance of the welfare losses reported. The figures presented in table 4.1 seem relatively small, but if we compare them to Lucas (2003 [65]) they are certainly not trivial. Using US data, Lucas calculates that the welfare cost of fluctuations in consumption around its trend is only about 0.05 of one percent of consumption. Although this model, and the utility function I use, differs from Lucas', this does give us some idea of

the significance of the effect of the behavioural biases in the model. Furthermore, there are two reasons to believe that this cost is understated.

- (1) Firstly, the idea of using future expectations from this model as the basis for welfare comparisons is clearly problematical. The fact that the model is not based on fully rational expectations means that ex-ante (or anticipated) welfare can be inconsistent with ex-post (or realised) welfare. Given that I use expectations to estimate the effect on welfare from period 41 onwards, there are grounds to believe that this measure of welfare underestimates the effect of the behavioural biases on actual welfare. Despite the fact that later periods are more heavily discounted, periods 41 to ∞ still account for around two thirds of inter-temporal welfare.
- (2) Secondly, as noted earlier, the main costs associated with asset price bubbles in the real world are as the result of credit rationing which often accompanies the bursting of a bubble. In this model, I have no financial market frictions which could cause such a credit crunch. The only way in which asset mis-pricings effect allocations in the real economy is via wealth effects. As Bean (2004 [15]) notes, "if the only macroeconomic consequences of booms and asset prices were via conventional wealth effects on aggregate demand, then they would constitute little more than a nuisance to monetary policy makers". Hence, the fact that I derive a significant cost to behavioural biases even in the absence of financial accelerator effects is noteworthy.

If my intention in this chapter were to try and provide a meaningful estimate of the cost of behavioural biases, then I would clearly need to address both of these issues. In fact, my intention is not to do that, but rather to assess the effectiveness of monetary policy in correcting for behavioural biases. In order to do this, I do not need an accurate estimate of the cost of behavioural biases, I only need for my estimate of that cost to be consistent across monetary policy regimes. For that reason, and given the computing power necessary to accurately measure ex-post welfare, and given the difficulty of introducing financial accelerator effects into a DSGE model, I will postpone the attempt to accurately estimate the costs of behavioural biases until further work.

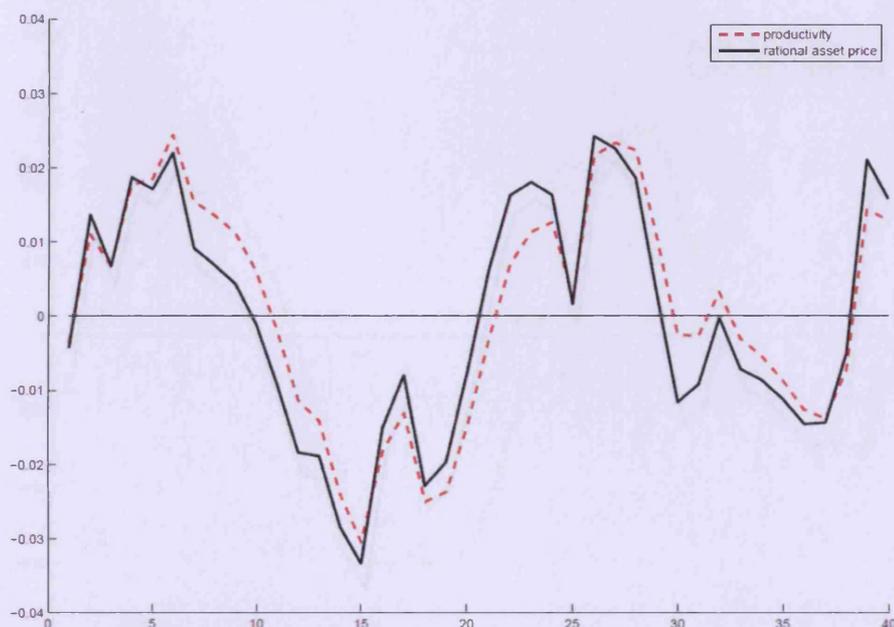
4.5.1. Intuition for the Result

In order to provide intuition for these results, I consider here a single simulation of the model.

Figure 4.7 shows the productivity process in this particular simulation. It also shows that the rational asset price mimics, almost exactly, the productivity process⁶. Figure 4.8 contrasts the asset price in the rational model with that in the behavioural model under a passive monetary regime, based on the same underlying

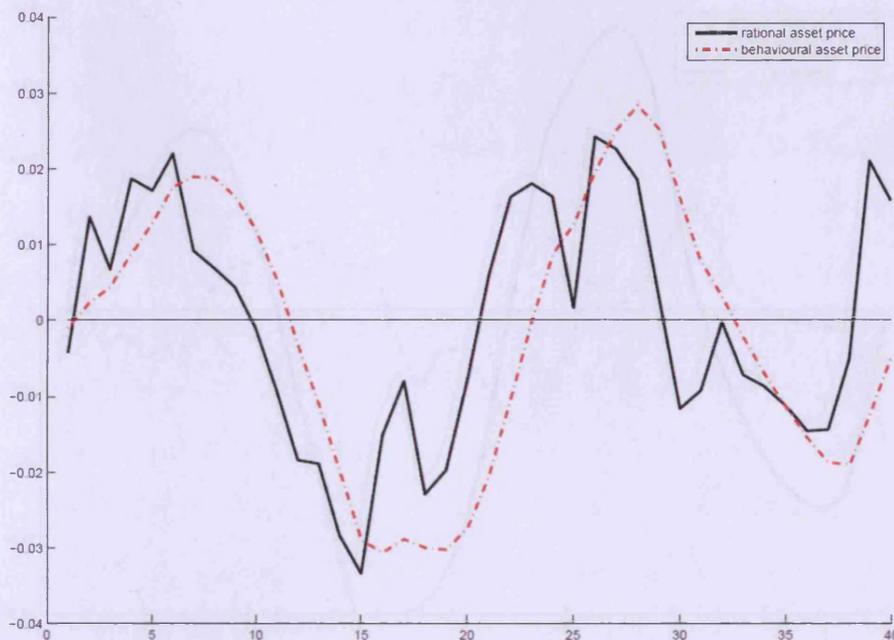
⁶The rational asset price illustrated in figure 4.7 is the asset price that I get when I simulate the model without any behavioural biases. This differs from the fundamental asset price under any simulation of the full model because the behavioural rules affect capital accumulation, and this in turn drives changes in what would then be a rational asset price.

Figure 4.7. The productivity process and rational asset price in a simulation of the model



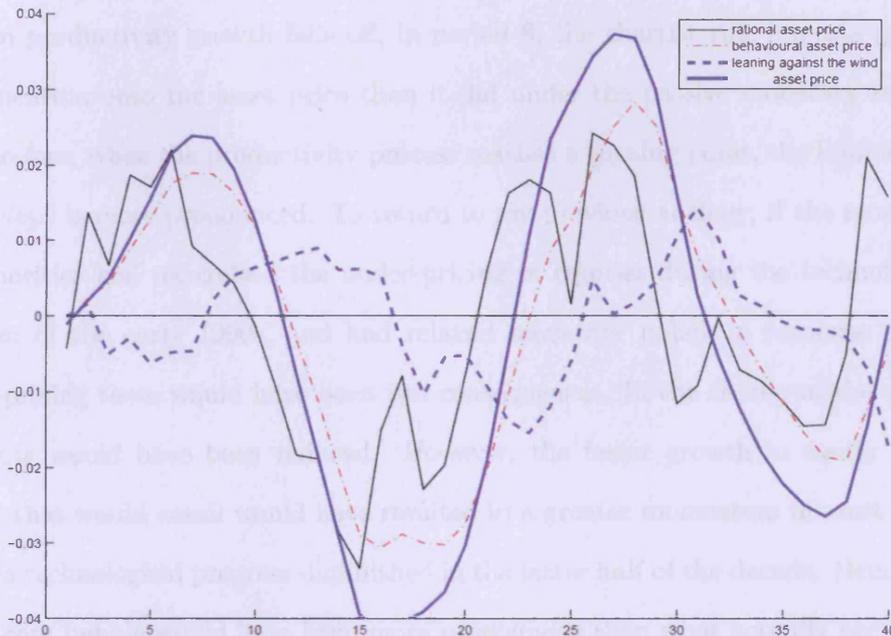
productivity process. We can see that the effect of the behavioural rules is that the asset price becomes slow to react to changes in productivity. This is because initially only fundamentalists react to the change in productivity. Chartists gradually jump on the bandwagon, but are slow to react when there is a turning point in productivity. This type of dynamics coincides with the type of story that is often told to explain asset price bubbles. For example, during the late 1980s and early 1990s there was a revolution in information technology, largely based on the growth of the internet, which did actually drive improvements in productivity. A

Figure 4.8. The asset price in the rational model and in the behavioural model with a passive monetary policy



plausible explanation of equity prices through the 1990s might argue that at first there was a slow response to the fundamental changes in productivity that were being driven by technological progress. However, once the response started to take place it accelerated at exactly the same time as the growth rate of productivity began to return to more normal levels. Hence, asset prices began to outstrip their fundamental value in what came to be known as the dot-com bubble. Similar stories have been told about the recent boom in house prices, but this time based on financial rather than informational innovations.

Figure 4.9. The asset price in the behavioural model with a ‘leaning against the wind’ monetary policy

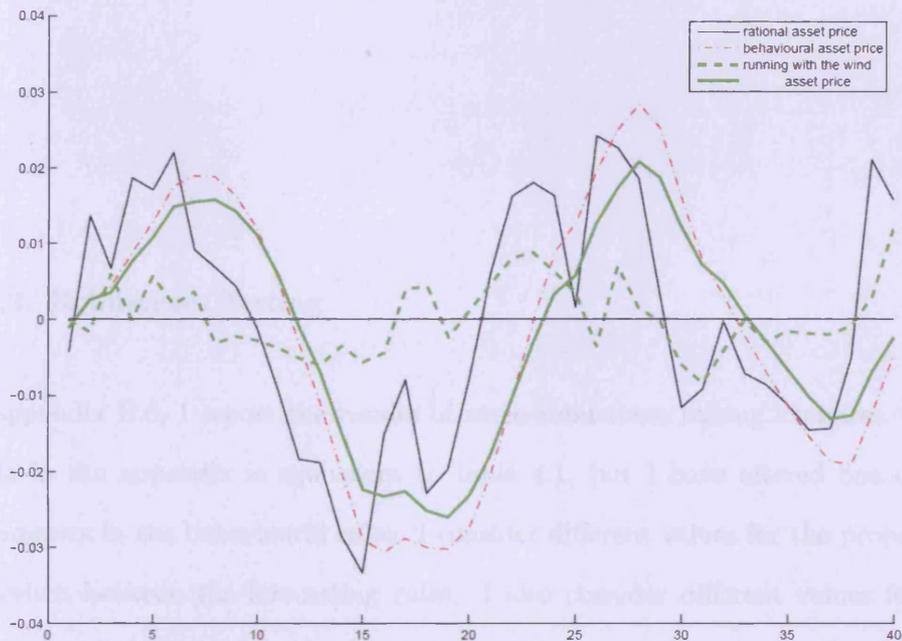


The question now is what effect monetary policy has in the model. Figure 4.9 illustrates the effect of a ‘leaning against the wind’ monetary policy, with $\zeta_Q = 0.5$. The dashed blue line is the monetary policy (interest rate) response to the asset mis-pricing. Monetary policy is expansionary (low interest rate) when the asset price was below fundamental in the previous period and it is contractionary (high interest rate) when the asset price was above fundamental in the previous period. With this monetary policy response, the asset price is described by the solid blue line. As we can see, in periods 1 to 7 of this simulation the monetary policy

response has the effect of driving the asset price closer to its fundamental value (approximated by the black line of the rational model). However, this means that when productivity growth falls off, in period 8, the chartist rule imparts greater momentum onto the asset price than it did under the passive monetary regime. Therefore, when the productivity process reaches a turning point, the bubble that develops is more pronounced. To return to my previous analogy, if the monetary authorities had recognised the under-pricing of equities during the technological boom of the early 1990s, and had relaxed monetary policy in response to the mis-pricing there would have been two consequences. In the short-run the under-pricing would have been reduced. However, the faster growth in equity prices that that would entail would have resulted in a greater momentum in asset prices when technological progress diminished in the latter half of the decade. Hence, the dot-com bubble would have been more pronounced than what actually occurred.

Figure 4.10 illustrates the effect of a ‘running with the wind’ monetary policy, with $\zeta_Q = -0.5$. The dashed green line is the interest rate response to the asset mis-pricing. Monetary policy is contractionary when the asset price was below fundamental in the previous period and it is expansionary when the asset price was above fundamental in the previous period. With this monetary policy response, the asset price is described by the solid green line. The effect here is the reverse of the ‘leaning against the wind’ policy. During periods 1 to 7 the asset price drifts further away from its fundamental value. However, when productivity growth

Figure 4.10. The asset price in the behavioural model with a 'running with the wind' monetary policy



begins to fall the momentum in the chartist forecast is less than under other policy specifications, and so the resulting bubble is less pronounced.

In this model bubbles appear when turning points in the productivity process occur. The crucial factor in how large those bubbles turn out to be is the rate of change of the asset price in the periods prior to any turning point. There is a momentum inherent in the chartist forecasting rule, and this momentum is greater when the rate of change in the asset price is higher. Since a 'leaning against

the wind' policy tends to promote rapid changes in asset prices, it also promotes greater momentum in the chartist rule, and more pronounced bubbles.

4.5.2. Robustness Testing

In appendix B.6, I report the results of some robustness testing exercises. Each table in the appendix is equivalent to table 4.1, but I have altered one of the parameters in the behavioural rules. I consider different values for the propensity to switch between the forecasting rules. I also consider different values for the uncertainty band in the fundamentalist rule, C .

The only case in which a 'leaning against the wind' policy increases welfare is when the switching parameter, v , is very large. Even in this case only a small response is beneficial, whilst an aggressive policy is extremely harmful.

4.6. Conclusions

The main result of the model that I have developed in this chapter is that ‘leaning against the wind’ monetary policies are counter-productive whilst ‘running with the wind’ policies can ameliorate the effects of behavioural biases. This result is clearly model specific. It relies on the particular characteristics of the behavioural forecasting rules on which the model is based. Although I have shown, in the previous chapter, that these simple heuristical rules can account for historical asset price dynamics, I do not want to claim that they are a realistic description of real world financial markets. I, therefore, need to be cautious in drawing policy conclusions from this model.

What is clear is that ‘leaning against the wind’ monetary policies cannot be relied upon to correct for behavioural biases, and in some cases will cause serious harm. It is not so clear that we should be supporting systematic use of ‘running with the wind’ policies. It is likely that this result is very model specific. The issue here is that the dynamics of the bubble, and the behavioural underpinnings of those dynamics, are extremely important in determining the most relevant policy response. The big problem, of course, is that the monetary authorities do not have much understanding of that behaviour.

It is important to note some caveats to what I have said so far. Firstly, my model ignores any financial accelerator effects that may exist. These would likely

exacerbate the costs of sub-rational behaviour. However, they will also serve to add further complexity to the picture.

Secondly, my model does not allow for any effect caused by expectations of future monetary policy. The model does not allow for the fact that a pre-commitment to a 'leaning against the wind' might cause private agents to hedge against mis-pricings by shunning chartist rules. However, we know that if there is an asset mis-pricing, it must be as the result of some irrationality in the market. If such irrationality exists then it is unclear as to why it would not prevent agents from responding optimally to policy pre-commitments. In Calvo-type models, monetary policy can correct for nominal rigidities by keeping inflation and expectations of inflation equal, through a commitment to an inflation target. In such models, rationality reigns and the monetary authorities are able to take advantage of this in predicting how private expectations will react to future monetary policy. As soon as we loosen the requirement for rationality, it becomes difficult for the monetary authorities to predict how policy affects private sector expectations.

Thirdly, it is important to note that in my analysis, I have only considered systematic monetary policy rules. It may well be possible that a one off monetary tightening or loosening could be effective in increasing welfare if the monetary authorities are able to identify turning points in the productivity process.

Bubbles can only exist in a world where agents are less than fully rational. Furthermore, the dynamics of bubbles are dependant upon the nature of that irrationality. The relevant policy response to a bubble depends upon the particular nature of the behavioural reaction of the market. We have only a very rudimentary understanding of such behaviours. What we do know suggests that such behaviours are complex and cause discontinuities in the relationships between different variables. For all these reasons, it is likely that any attempt by central banks to try to influence market psychology is likely to have highly unpredictable outcomes.

In conclusion, I am arguing that a systematic monetary policy response to asset mis-pricings is unlikely to enhance welfare. The reason for this is that monetary authorities lack the full information necessary to implement an effective policy.

CHAPTER 5

Overall Conclusions

Rationality can be viewed as a normative model of decision making. It defines the basis on which we should make decisions. Positive models of decision making, on the other hand, seek to describe how people actually make decisions in the real world. Behavioural theories belong to this class. What is clear is that the standard model of rationality, as expressed by subjective expected utility theory, fails to provide a complete explanation of how people in the real world make decisions. It is a moot point as to whether this is because there is some lacking in our axiomatisation of rationality, or whether people simply fail to live up to the rational ideal.

A useful distinction to draw is between those models which are prefaced on a normative account of human decision making, and those that are based on a positive account. We are used, in some fields of economics, to distinguishing between models which have explanatory and conceptual powers, on the one hand, and those which have descriptive powers, on the other. For example, nobody believes that perfectly competitive markets actually exist in the real world. The model of perfect competition is based on many assumptions, including those of perfect information and homogenous goods, which simply are not observed in the real world. The

model of perfect competition, however, provides us with useful concepts against which to compare the real world; an understanding of some of the relationships between different economic variables; and an ideal towards which economic policy might focus. The theory of perfect competition provides us with an account of how an ideal world might be, not how the real world actually is. This type of model can enhance our understanding of how the economy works just as much as more empirically accurate ones can.

A similar dichotomy can be observed in macroeconomics. Real business cycle theory (and classical theories more generally) can be seen as a counterpart to perfect competition. It describes how an economy 'should' function, if only people were rational, there were no informational constraints, and no-one could exert market power. Such theory tells us a lot about the relationships between different macroeconomic variables and, arguably, describes how an ideal economy might function. However, unsurprisingly, it does very poorly in describing real world economies. There is a well established consensus that, in order to build macroeconomic models that can actually account for real world experiences, we must include some form of nominal rigidity. Clearly, nominal rigidities must arise out of some imperfection when compared to the real business cycle model. This might either be due to imperfect information (as is popular in neoclassical accounts), or due to some type of menu cost, or it may be due to some sub-optimal behaviour by agents. This final explanation is referred to as 'money illusion'.

Even though introducing nominal rigidities tends to improve the performance of macroeconomic models in replicating the real world, New Keynesian models are still poor at doing this in many respects. In particular, they struggle to recreate the volatility and kurtosis in equity returns.

In chapter 3, I established the fact that simple behavioural rules can help us to explain the dynamics of asset prices. There are, of course, other explanations of why asset prices behave in the way that they do, but a growing consensus is emerging that human behaviour is an essential ingredient in any comprehensive explanation. The significance of my results in chapter 3 is that I am able to explain all of the dynamics of the FTSE all-share index using a behavioural model. What's more, I am able to do so with a relatively simple set of heuristical rules.

I do not want to claim that the heuristical rules in the model in chapter 3 provide an accurate account of the way in which agents in financial markets actually behave. Rather, my claim is that these equations are a construct that capture some of the dynamics that are caused by irrationality in those markets. In much the same way, no one believes that Calvo contracts are a description of any phenomenon that actually exists in the real world. Instead, they are a construct which encapsulate the effects of an imperfection in real world economies (whether that be money illusion, menu costs, imperfect information or a combination of all three).

I can persist with this analogy with nominal rigidities and Calvo contracts in explaining the model that I put forward in chapter 4. In the same way as I include Calvo contracts in the model as a proxy for nominal rigidity, I include

the behavioural equations as a proxy for sub-rational decision making in financial markets.

I began this thesis by posing the question of whether considerations of the ways in which people deviate from rationality are important in explaining the behaviour of the macroeconomy. If so, the next obvious question is what ramifications this may hold for macroeconomic policy formulation. Following the previous chapters, I am now hopefully in a reasonable position to lend a fresh perspective to these questions.

My conclusion from chapter 3 is that behavioural considerations can indeed help us to explain asset price dynamics. In turn, this might help us to explain further macroeconomic variables. However, it does not hold that the existence of noise trader risk gives rise to an additional target for monetary policy.

By its very nature, irrationality is complex. Kahneman and Tversky highlight three different heuristics that give rise to innumerable biases. Which of these biases are most salient and what happens when different biases are in conflict with each other is unclear. Even their axiomatisation of these concepts as prospect theory allows for ambiguities in the outcomes of very similar decisions. These ambiguities are referred to as framing effects.

As we discovered in chapter 3, even when we begin with a well specified set of simple heuristical rules, discontinuities in the relationships between behavioural variables leads to complex and even chaotic outcomes. In this particular model, herding effects are apparent as an example of this complexity.

Indeed, I find in chapter 4 that in that particular model, under most parameterisations, a counter-intuitive ‘running with the wind’ monetary policy is effective in increasing welfare. This result arises because such a policy tends to prevent momentum in the chartist rule from developing quickly. Clearly, for a central bank to be able to anticipate such an effect, it would have to have an intimate understanding of the behavioural norms that exist in the market. It is unrealistic to posit such a level of understanding.

It is appropriate to close this thesis with reference to the recent periods of financial turmoil, notably the crash of the dot-com bubble in 2001 and the more recent ‘credit crunch’.

Few commentators have attempted to explain the dot-com equities bubble or the housing bubble that preceded the credit crunch without some reference to the role of sub-rational human decision making. The dynamics of both of these bubbles are consistent with the ideas inherent in the behavioural models that I have

explored in this thesis. Both bubbles have been described as the result of momentum that built up when there were improvements in the underlying fundamentals, but which was not checked when the improvements in fundamentals abated.

If these bubbles were indeed the result of chartist speculation, as exemplified in my models, then the results of chapter 4 concur with Alan Greenspan's assessment that:

"the notion that a well timed incremental tightening [of monetary policy] could have been calibrated to prevent the ... bubble is almost surely an illusion" (Greenspan 2002 [49])

The reason that I reach this conclusion, however, is very different to Greenspan's explanation. He argues that if central banks had access to information that would allow them to reasonably calculate the fundamental value of assets, then private agents would be able to do so as well and bubbles would be very unlikely to develop. I do not contest the fact that assessing the fundamental value of assets is problematic, though Cecchetti et al (2000 [27]) provide strong arguments that misalignments are often identifiable. However, what my model illustrates is that, even if central banks can perfectly measure the fundamental asset price, there are reasons why it would be imprudent for them to use monetary policy to try to correct for a mis-pricing.

For a bubble to exist in the first place there must be some imperfection in the asset market. At least some proportion of the agents in the market must have

beliefs that are incompatible with rational expectations. Furthermore, there must exist some barriers to arbitrage which prevent those who have rational beliefs from driving the price back to fundamental. It would take a brave individual to predict what effects monetary intervention would have in such a market. What I can say is that, under the particular behavioural rules examined in chapter 4, a 'leaning against the wind' monetary policy would be detrimental.

Seeing as most of the welfare costs associated with asset mis-pricing arise from financial accelerator effects, the authorities would be best advised to aim their efforts at lessening financial frictions and reducing leverage. The analysis in this thesis suggests that a prudent central bank should not attempt to mitigate these costs by influencing market psychology.

References

- [1] Abel, A. (1990). "Asset prices under habit formation and catching up with the Joneses", *American Economic Review* 80, pages 38-42.
- [2] Adams, J.S. (1963). "Toward an understanding of inequity", *Journal of Abnormal and Social Psychology* 67, pages 422-436.
- [3] Akerlof, G.A. (1982). "Labor Contracts as Partial Gift Exchange", *Quarterly Journal of Economics* 97, pages 543-69.
- [4] Akerlof, G.A. (2002). "Behavioral Macroeconomics and Macroeconomic Behavior", *American Economic Review*, vol. 92(3), pages 411-433.
- [5] Akerlof, G.A. and Yellen, J.L. (1985). "A Near-rational Model of the Business Cycle, with Wage and Price Intertia", *Quarterly Journal of Economics* 100(5), pages 823-38.
- [6] Akerlof, G.A. and Yellen, J.L. (1985). "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?", *American Economic Review* 75(4), pages 708-20.
- [7] Allais, M. (1953). "Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole Americaine", *Econometrica* 21(4), pages 503-546.
- [8] Aristotle. *Nicomachean Ethics*. London: Penguin Books (2004).
- [9] Assenmacher-Wesche, K. and Gerlach, S. (2008). "Ensuring Financial Stability: Financial Structure and the Impact of Monetary Policy on Asset Prices", *CEPR Discussion Papers* 6773.

- [10] Bachelier, L. (1900/1964). "Theorie de la speculation". In: Cootner, P.H. ed. *The random character of stock market prices*. Cambridge, Massachusetts: MIT Press, pages 17-78.
- [11] Banz, R.W. (1981). "The Relationship Between Return and Market Value of Common Stocks", *Journal of Financial Economics* 9(1), pages 3-18.
- [12] Barberis, N. et al. (2001). "Prospect theory and asset prices", *Quarterly Journal of Economics* 116(1), pages 1-53.
- [13] Barberis, N., Schleifer, A. and Vishny, R. (1998). "A model of investor sentiment", *Journal of Financial Economics* 49(3), pages 307-343.
- [14] Barberis, N. and Thaler, R. (2002). "A survey of behavioural finance". In: Constantinides, G. et al. eds. *Handbook of the Economics of Finance*. JAI Press.
- [15] Bean, C. R. (2004). "Asset Prices, Financial Instability, and Monetary Policy", *American Economic Review* 94(2), pages 14-18.
- [16] Bernanke, B. and Gertler, M. (2000). "Monetary Policy and Asset Price Volatility". *NBER Working Papers* 7559, National Bureau of Economic Research, Inc.
- [17] Bernanke, B. and Gertler, M. (2001). "Should Central Banks Respond to Movements in Asset Prices?" *American Economic Review*, 91(2), pages 253-257.
- [18] Bernanke, B., Gertler, M. and Gilchrist, S. (1999). "The financial accelerator in a quantitative business cycle framework", in: J. B. Taylor & M. Woodford (ed.), *Handbook of Macroeconomics*, 1(1), pages 1341-1393. Elsevier.
- [19] Benartzi, S. and Thaler, R.H. (1995). "Myopic loss aversion and the equity premium puzzle", *Quarterly Journal of Economics* 110(1), pages 73-92.
- [20] Bernoulli, D. (1738/1954). "Exposition of a new theory on the measurement of risk", *Econometrica* 22(1), pages 23-26.
- [21] Bernoulli, J. (1713). *Ars conjectandi*. Basel.

- [22] Birnbaum, M. (1983). "Base rates in Bayesian inference: Signal detection analysis of the cab problem", *American Journal of Psychology* 96, pages 85-94.
- [23] Bordo, M.D. and Jeanne, O. (2002). "Monetary Policy and Asset Prices: Does "Benign Neglect" Make Sense?", *IMF Working Papers* 02/225, International Monetary Fund.
- [24] Brock, Lakonishok and LeBaron. (1992). "Simple technical trading rules and the stochastic properties of stock returns", *Journal of Finance* 47, pages 1731-1764.
- [25] Campbell, J.Y. and Cochrane, J.H. (1999). "Force of habit: A consumption-based explanation of aggregate stock market behavior", *Journal of Political Economy* 107(2), page 205-251.
- [26] Canzoneri, M.B., Cumby, R.E. and Diba, B.T. (2007). "The Cost of Nominal Inertia in NNS Models", *Journal of Money, Credit and Banking* 39(7), pages 1563-1586.
- [27] Cecchetti, S., Genberg, H., Lipsky, J. and Wadhvani, S. (2000). "Asset Prices and Central Bank Policy", *Geneva Reports on the World Economy* 2, International Centre for Monetary and Banking Studies and Centre for Economic Policy Research.
- [28] Cosmides, L. & Tooby, J. (1996). "Are humans good intuitive statisticians after all?: Rethinking some conclusions of the literature on judgment under uncertainty", *Cognition* 58, pages 1-73.
- [29] Cowles, A. and Jones, H. (1937). "Some a posteriori probabilities in stock market action", *Econometrica* 5, pages 280-294.
- [30] Daniel, K., Hirshleifer, D. and Subrahmanyam, A. (1998). "Investor psychology and security market under- and overreactions", *Journal of Finance* 53(6), pages 1839-1885.
- [31] De Grauwe, P. and Grimaldi, M. (2006). *The exchange rate in a behavioral finance framework*. Princeton University Press.

- [32] De Long, J.B. et al. (1990). "Noise Trader Risk in Financial Markets", *Journal of Political Economy* 98(4), pages 703-738.
- [33] Eaton, B.C. and White, W. (1982). "Agent Compensation and the Limits of Bonding", *Economic Inquiry* 20, pages 330-43.
- [34] Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms", *Quarterly Journal of Economics* 75(4), pages 643-669.
- [35] Epstein, L.G. and Zin, S.E.. (1991). "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis", *Journal of Political Economy* 99, pages 263-286.
- [36] Fama, E.F. (1970). "Efficient capital markets: A review of theory and empirical work", *Journal of Finance* 25(2), pages 383-417.
- [37] Fama, E.F. and French, K.R. (1988). "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics* 22(1), pages 3-25.
- [38] Fehr, E. and Gächter, S. (2000). "Fairness and Retaliation: The Economics of Reciprocity", *Journal of Economic Perspectives*. 14(3), pages 159-181.
- [39] Fehr, E. and Tyran, J.R. (2001). "Does Money Illusion Matter?", *American Economic Review* 91, pages 1239-1262.
- [40] Fiedler, K. (1988). "The dependence of the conjunction fallacy on subtle linguistic factors", *Psychological Research* 50, pages 123-129.
- [41] Fisher, I. (1928). *The Money Illusion*. New York: Adelphi.
- [42] Frankel, J.A. and Froot, K.A. (1986). "The Dollar as a Speculative Bubble: A Tale of Fundamentalists and Chartists", NBER Working Papers 1854.
- [43] Frankel, J.A. and Froot, K.A. (1990). "Chartists, Fundamentalists, and Trading in the Foreign Exchange Market", *American Economic Review* 80(2), pages 181-85.
- [44] French, K.R. (1980). "Stock Returns and the Weekend Effect", *Journal of Financial Economics* 8(1), pages 55-69.

- [45] Froot, K.A. and Dabora, E.M. (1999). "How are stock prices affected by the location of trade?", *Journal of Financial Economics* 53(2), pages 189-216.
- [46] Gigerenzer, G. (1991). "How to make cognitive illusions disappear: Beyond heuristics and biases", *European Review of Social Psychology* 2, pages 83-115.
- [47] Gigerenzer, G. et al. (1988). "Presentation and content: The use of base rates as a continuous variable", *Journal of experimental psychology: Human perception and performance* 14(3), pages 513-525.
- [48] Greenspan, A. (1995). "The Challenge of Central Banking in a Democratic Society", *Annual Dinner of The American Enterprise Institute for Public Policy Research*, Washington, D.C. <http://www.federalreserve.gov/boarddocs/speeches/1996/19961205.htm>
- [49] Greenspan, A. (2002). "Opening Remarks", *Rethinking Stabilisation Policy*, Federal Reserve Bank of Kansas City.
- [50] Gruen, D., Plumb, M. and Stone, A. (2005). "How Should Monetary Policy Respond to Asset-Price Bubbles?", *International Journal of Central Banking* 3(1), pages 1-33.
- [51] Hacking, I. (1980). "Strange expectations", *Philosophy of Science* 47(4), pages 562-567.
- [52] Harris, L. (1986). "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns", *Journal of Financial Economics* 16(1), page 99.
- [53] Honderich, E. (2005). *The Oxford Companion to Philosophy*. Oxford University Press.
- [54] Kahneman, D. and Tversky, A. (1972). "Subjective probability: A judgement of representativeness", *Cognitive Psychology*, pages 430-454.
- [55] Kahneman, D. and Tversky, A. (1973). "On the psychology of prediction", *Psychological Review* 80, pages 237-251.
- [56] Kahneman, D. and Tversky, A. (1979). "Prospect Theory: an Analysis of Decision Under Risk", *Econometrica* 47(2), pages 263-291.

- [57] Kahneman, D. et al. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.
- [58] Kay, J. (2004). *The truth about markets*. London: Penguin.
- [59] Kendall, M. (1953). "The analysis of economic time series", *Journal of the Royal Statistical Society* 96, pages 11-25.
- [60] Keynes, J.M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.
- [61] Klayman, J. and Ha, Y.-W. (1987). "Confirmation, disconfirmation and information in hypothesis testing", *Psychological Review* 94(2), pages 211-228.
- [62] Lo, A.W. and Mackinlay, A.C. (1988). "Stock market prices do not follow random walks: Evidence from a simple specification test", *Review of Financial Studies* 1(1), pages 41-66.
- [63] Loewenstein, G. and Prelec, D. (1992). "Anomalies in intertemporal choice: Evidence and an interpretation", *Quarterly Journal of Economics* 107, pages 573-598.
- [64] Lucas, R.E.J. (1978). "Asset Prices in an Exchange Economy", *Econometrica* 46(6), pages 1429-1445.
- [65] Lucas, R.E.J. (2003). "Macroeconomic Priorities", *American Economic Review*, 93(1), pages 1-14.
- [66] Machina, M.J. (1987). "Choice under uncertainty: Problems solved and unsolved", *Journal of Economic Perspectives* 1(1), pages 121-154.
- [67] Markowitz, H.M. (1952). "The Utility of Wealth", *Journal of Political Economy* 60(2), pages 151-158.
- [68] Meenagh, D., Minford, P. and Peel, D (2007). "Simulating stock returns under switching regimes - a new test of market efficiency", *Economics Letters* 94, pages 235-239.
- [69] Mehra, R. and Prescott, E.C. (1985). "The Equity Premium: A Puzzle", *Journal of Monetary Economics* 15(2), pages 145-161.

- [70] Minford, P., Theodoridis, K. and Meenagh, D. (2007). "A bootstrap test for the dynamic performance of DSGE models - an outline and some experiments", *Cardiff Economics Working Papers* E2007/2, Cardiff University.
- [71] Mishkin, F.S. (2007). "Housing and the monetary transmission mechanism", *Finance and Economics Discussion Series* 2007-40, Board of Governors of the Federal Reserve System (U.S.).
- [72] Oaksford, M. and Chater, N. (1994). "A rational analysis of the selection task as optimal data selection", *Psychological Review* 101(4), pages 608-631.
- [73] Oaksford, M. and Chater, N. (1996). "Rational explanation of the selection task", *Psychological Review* 103(2), pages 381-391.
- [74] Pascal, B. (1660/1961). *Pensees*. Penguin Books.
- [75] Rabin, M. (2000). "Risk Aversion and Expected-Utility Theory: A Calibration Theorem", *Econometrica* 68(5), pages 1281-1292.
- [76] Rietz, T.A. (1988). "The equity risk premium: a solution", *Journal of Monetary Economics* 22, pages 117-131.
- [77] Ritter, J.R. (1991). "The Long-Run Performance of Initial Public Offerings", *Journal of Finance* 46(1), pages 3-27.
- [78] Roberts, H. (1967). *Statistical versus clinical prediction of the stock market*. University of Chicago.
- [79] Roy, D. (1952). "Quota restriction and goldbricking in a machine shop", *American Journal of Sociology* 67(2), pages 427-42.
- [80] Rozeff, M. and Kinney, W. (1976). "Capital market seasonality: The case of stock returns", *Journal of Financial Economics* 3, pages 370-402.
- [81] Samuelson, P.A. (1965). "Proof that properly anticipated prices fluctuate randomly", *Industrial Management Review* 6(2), pages 41-49.
- [82] Savage, L.J. (1954). *The foundations of statistics*. New York: Wiley.
- [83] Shafir, E., Diamond, P.A., and Tversky, A. (1997). "On Money Illusion", *Quarterly Journal of Economics* 112, pages 341-74.

- [84] Shapiro, C. and Stiglitz, J. (1984). "Equilibrium Unemployment as a Worker Discipline Device", *American Economic Review* 74(3), pages 433-444.
- [85] Simon, H.A. (1956). "Rational choice and the structure of the environment", *Psychological Review* 63, pages 129-138.
- [86] Simon, H.A. (1987). "Bounded rationality". In: Eatwell, J. et al. eds. *The new Palgrave: A dictionary of economics*. Macmillan.
- [87] Stiglitz, J. (1974). "Wage Determination and Unemployment in L.D.C.'s: The Labor Turnover Model", *Quarterly Journal of Economics* 88, pages 194-227.
- [88] Stiglitz, J. (1976). "Prices and Queues as Screening Devices in Competitive Markets", *IMSSS Technical Report* No. 212, Stanford University. [reprinted in *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, D. Gale and O. Hart (eds.), Cambridge:MIT Press, 1992 pages 128-166]
- [89] Tversky, A. and Kahneman, D. (1973). "Availability: A heuristic for judging frequency and probability", *Cognitive Psychology*, pages 207-232.
- [90] Thaler, R.H. and Sunstein, R.S. (2008). *Nudge*. Yale University Press.
- [91] Tversky, A. and Kahneman, D. (1974). "Judgment under uncertainty: Heuristics and biases", *Science* 185, pages 1124-1131.
- [92] Tversky, A. and Kahneman, D. (1982a). "Evidential impact of base rates". In: Kahneman, D. et al. eds. *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press, pages 153-160.
- [93] Tversky, A. and Kahneman, D. (1982b). "Judgements of and by representativeness". In: Kahneman, D. et al. eds. *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press, pages 84-98.
- [94] Tversky, A. and Kahneman, D. (1992). "Advances in prospect theory: Cumulative representation of uncertainty", *Journal of Risk and Uncertainty* 5, pages 297-323.

- [95] von Neumann, J. and Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- [96] Wason, P.C. (1960). "On the failure to eliminate hypotheses in a conceptual task", *Quarterly Journal of Experimental Psychology* 12, pages 129-140.
- [97] Wason, P.C. (1966). "Reasoning". In: Foss, B.M. ed. *New horizons in psychology*. Harmondsworth: Penguin, pages 135-161.
- [98] Wason, P.C. (1968). "Reasoning about a rule", *Quarterly Journal of Experimental Psychology* 20, pages 273-281.
- [99] Working, H. (1934). "A random difference series for use in the analysis of time series", *Journal of the American Statistical Association* 29, pages 11-24.
- [100] Yellen, J.L. (1984). "Efficiency Wage Models of Unemployment", *American Economic Review* 74(2), pages 200-205.

APPENDIX A

A Behavioural Asset Price Model

A.1. Fortran code

```
PROGRAM C5grid

!Purpose: end of column 70-->
! To search for the best fit parameters for the subroutine C5model
!
!Record of revisions:
! Date Programmer Description of change
! ==== =====
! 26/09/07 Rhys ap Gwilym Original code

IMPLICIT NONE

!Declare variables:
INTEGER, PARAMETER :: double=8 !double precision
INTEGER :: i,j,k,l,counter,inseed !loop and counter variables
REAL, DIMENSION(10) :: parameters !model parameters
REAL, DIMENSION(6) :: startpars, bestpars !starting/best model parameters
REAL, DIMENSION(6) :: maxpar, minpar, step, stepsize !max and min parameter
                                                    !values, step size

REAL, DIMENSION(7,13) :: x !
EXTERNAL G05FDF, G05CBF !pseudo-random number generator
REAL(KIND=double) :: crit, bestcrit !(best) critical value
INTEGER, PARAMETER :: ss=10 !number of stochastic simulations
                        ! at each parameterisation

!Fixed parameters
parameters(5)=1.0 !mu
parameters(6)=0.0 !phi
parameters(7)=0.0 !r
parameters(9)=0.0 !p0

!Parameters to search over:
```

```

! psi beta rho gamma c sdelta
startpars=(/ 0.525,0.965,0.4, 3.75, 0.425,0.04 /)
step= (/ 0.1, 0.1, 0.1, 1.0, 0.1, 0.01 /)
minpar= (/ 0.3, 0.5, 0.2, 1.0, 0.30, 0.01 /)
maxpar= (/ 0.7, 0.95, 0.6, 5.0, 0.50, 0.05 /)

!Calculate critical value for starting parameters
parameters(1)=startpars(1)
parameters(2)=startpars(2)
parameters(3)=startpars(3)
parameters(4)=startpars(4)
parameters(8)=startpars(5)
parameters(10)=startpars(6)
C CALL C5 (parameters, inseed, ss, crit)

bestpars=startpars
bestcrit=100000

inseed=0

steps: DO l=0,2
DO i=1,6
stepsize(i)=step(i)/(2.0**l)
END DO

parsearch: DO

inseed=inseed+100

!put best parameters and critical value in first row of x matrix
x=0.0
DO j=1,6
DO k=1,13
x(j,k)=bestpars(j)
END DO
END DO
C x(7,1)=bestcrit

!vary one parameter in every other row of x matrix
DO i=1,6
DO j=1,2

```

```

k=(i*2)+j-1
x(i,k)=x(i,k)+(2*j-3)*stepsize(i)
IF (x(i,k) > maxpar(i)) x(i,k)=maxpar(i)
IF (x(i,k) < minpar(i)) x(i,k)=minpar(i)
END DO
END DO

!calculate critical values for all rows of x matrix
DO k=1,13
parameters(1)=x(1,k)
parameters(2)=x(2,k)
parameters(3)=x(3,k)
parameters(4)=x(4,k)
parameters(8)=x(5,k)
parameters(10)=x(6,k)
CALL C5 (parameters, inseed, ss, crit)
x(7,k)=crit
END DO

WRITE(1,*) inseed
WRITE(1,'(7F10.4)') x

!search for lowest critical value and transfer those parameters to bestpars
bestcrit=x(7,1)
counter=1
DO k=2,13
IF (x(7,k) < bestcrit) THEN
counter=k
bestcrit=x(7,k)
DO i=1,6
bestpars(i)=x(i,k)
END DO
END IF
END DO

!if no improvement in critical value, reduce step size
IF (counter==1) EXIT parsearch

END DO parsearch
END DO steps

END PROGRAM

```

```

SUBROUTINE C5 (parameters, inseed, ss, crit)

!Purpose: end of column 70-->
! To simulate the model ss number of times given a particular
! parameterisation and return a critical value
!
!Record of revisions:
! Date Programmer Description of change
! ==== =====
! 06/07/07 Rhys ap Gwilym Original code

IMPLICIT NONE

INTEGER, PARAMETER :: double=8 !double precision

!Declare calling parameters:
REAL, INTENT(IN), DIMENSION(10) :: parameters !model parameters
INTEGER, INTENT(IN) :: inseed !input seed
INTEGER, INTENT(IN) :: ss !number of stochastic simulations at each
! parameterisation
REAL(KIND=double), INTENT(OUT) :: crit !critical value, store old parameter

!Declare local variables:
INTEGER :: i,j !loop and counter variables
INTEGER, PARAMETER :: t=1000 !number of simulation periods
REAL(KIND=double), DIMENSION(t+2) :: p, pstar, profits !vectors of actual and
! fundamental prices and profits
REAL(KIND=double), DIMENSION(t) :: epsilon !vectors of fundamental
! and price shocks
EXTERNAL G05FDF, G05CBF !pseudo-random number generator
REAL(KIND=double), DIMENSION(t+2) :: wf, wc, epf, epc !weights, expected
! prices
REAL(KIND=double), DIMENSION(t+2) :: pif, pic, varf, varc !profits, risks
REAL(KIND=double), DIMENSION(175) :: dlnp !D(ln[generatedFTSEprice])
REAL(KIND=double), DIMENSION(4) :: momdlnp, avemom !(average) moments of
! price series

REAL(KIND=double) :: crit1 !
INTEGER :: seed !seed for random number generators

crit=0.0

ssloop: DO j=1,ss

```

```

seed = inseed + j + 25

!Generate pstar
CALL genpstar (t, seed, pstar, profits)

!Run model
CALL C3model (t, parameters, seed, pstar, p, wf, wc,
> epf, epc, pif, pic, varf, varc)

!Calculate moments of the generated FTSE series (final 175 observations of
! dlnp)
DO i=1,175
dlnp(i) = p(t+2-175+i) - p(t+2-175+i-1)
END DO

CALL moments (175, dlnp, momdlnp)

!Calculate critical value
crit1 = (((momdlnp(1) - 0.004593708026)/0.004593708026) ** 2)
> + (((momdlnp(2) - 0.010507946551)/0.010507946551) ** 2)
> + (((momdlnp(3) - 0.017838538603)/0.017838538603) ** 2)
> + (((momdlnp(4) - 6.571520857831)/6.571520857831) ** 2)

crit = crit + crit1

END DO ssloop

crit = crit / ss

C avemom(2) = avemom(2) / ss
C avemom(3) = avemom(3) / ss
C avemom(4) = avemom(4) / ss

C !Calculate critical value
C crit = (((avemom(1) - 0.004593708026)/0.004593708026) ** 2)
C > + (((avemom(2) - 0.010507946551)/0.010507946551) ** 2)
C > + (((avemom(3) - 0.017838538603)/0.017838538603) ** 2)
C > + (((avemom(4) - 6.571520857831)/6.571520857831) ** 2)

RETURN
END SUBROUTINE

```

```

SUBROUTINE C3model (t, parameters, seed, pstar, p, wf, wc,
> epf, epc, pif, pic, varf, varc)

!Purpose: end of column 70-->
! To numerically simulate the more complex De Grauwe type model
!
!Record of revisions:
! Date Programmer Description of change
! ==== =====
! 12/03/07 Rhys ap Gwilym Original code
! 15/03/07 Rhys ap Gwilym Add in price setting shock
! 20/03/07 Rhys ap Gwilym Add history to chartists' forecast
! 22/03/07 Rhys ap Gwilym Import parameters and errors from external files
! 27/03/07 Rhys ap Gwilym Update to chapter 3 model
! 09/05/07 Rhys ap Gwilym Convert to subroutine

IMPLICIT NONE

INTEGER, PARAMETER :: double=8 !double precision

!Declare calling parameters:
INTEGER, INTENT(IN) :: t !number of simulation periods
REAL, INTENT(IN), DIMENSION(10) :: parameters !model parameters
INTEGER, INTENT(IN) :: seed !seed for random number generators
REAL(KIND=double), INTENT(IN), DIMENSION(t+2) :: pstar !vector of fundamental
! price
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: p !actual price
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: wf !fundamentalist weight
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: wc !chartist weight
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: epf !fundamentalist
! expectation
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: epc !chartist expectation
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: pif !fundamentalist profit
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: pic !chartist profit
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: varf !fundamentalist risk
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: varc !chartist risk

!Declare local variables:
INTEGER :: i,j !loop variables
INTEGER, DIMENSION(t+2) :: tp !time period
REAL :: psi !speed of adjustment to fundamental

```

```

REAL :: beta !chartists' extrapolation
REAL :: rho !chartists' memory
REAL :: gamma !rate of revision of forecasting rules
REAL :: mu !coefficient of risk aversion
REAL :: phi !risk perception
REAL :: r !interest rate
REAL :: c !transaction cost
REAL :: p0 !initial condition
REAL(KIND=double) :: sdeta !standard deviation of price shocks
REAL(KIND=double), DIMENSION(t+2) :: eta !pricing error
REAL(KIND=double), DIMENSION(t+2) :: s, muf !supply of asset, time varying
                                     ! risk aversion
EXTERNAL G05FDF, G05CBF !pseudo-random number generator

```

```

!Initialisation values

```

```

psi = parameters(1)
beta = parameters(2)
rho = parameters(3)
gamma = parameters(4)
mu = parameters(5)
phi = parameters(6)
r = parameters(7)
c = parameters(8)
p0 = parameters(9)
sdeta = parameters(10)
tp(1) = -1
DO i = 1,2
p(i) = 0.0D0
wf(i) = 0.5D0
wc(i) = 0.5D0
epf(i) = 0.0D0
epc(i) = 0.0D0
pif(i) = 0.0D0
pic(i) = 0.0D0
varf(i) = 0.0D0
varc(i) = 0.0D0
tp(i+1) = tp(i) + 1
END DO
s = 0
p(1) = pstar(1)
p(2) = pstar(2) + p0
epf(1) = pstar(1)
epf(2) = pstar(2)
epc(1) = pstar(1)

```

```
epc(2) = pstar(2)
```

```
!Generate pricing errors
CALL G05CBF(2*seed)
CALL G05FDF(0.0D0,sdeta,t,eta)
```

```
loop1: DO i = 3,t+2
```

```
!define time periods
tp(i) = tp(i-1) + 1
```

```
!Calculate risks:
varf(i) = (rho * varf(i-1)) + ((1 - rho) *
> ((epf(i-2) - p(i-1))**2))
varc(i) = (rho * varc(i-1)) + ((1 - rho) *
> ((epc(i-2) - p(i-1))**2))
```

```
!Calculate fundamentalist profit:
IF (epf(i-1) > ((1 + r) * p(i-1))) THEN
pif(i) = (p(i-1) - ((1 + r) * p(i-2)))
ELSE IF (epf(i-1) < ((1 + r) * p(i-1))) THEN
pif(i) = -(p(i-1) - ((1 + r) * p(i-2)))
ELSE IF (epf(i-1) == ((1 + r) * p(i-1))) THEN
pif(i) = 0
END IF
```

```
!Calculate chartist profit:
IF (epc(i-1) > ((1 + r) * p(i-1))) THEN
pic(i) = (p(i-1) - ((1 + r) * p(i-2)))
ELSE IF (epc(i-1) < ((1 + r) * p(i-1))) THEN
pic(i) = -(p(i-1) - ((1 + r) * p(i-2)))
ELSE IF (epc(i-1) == ((1 + r) * p(i-1))) THEN
pic(i) = 0
END IF
```

```
!Calculate fundamentalist and chartist weights:
muf(i) = mu / (1 + (phi * ABS(p(i-1) - pstar(i-1))))
wfi(i) = EXP(gamma * (pif(i) - (muf(i) * varf(i)))) /
> (EXP(gamma * (pif(i) - (muf(i) * varf(i)))) +
```

```

> EXP(gamma * (pic(i) - (mu * varc(i))))
wc(i) = 1 - wf(i)

!Calculate fundamentalist price expectation:
IF (ABS(p(i-1) - pstar(i-1)) > c) THEN
epf(i) = p(i-1) - (psi * (p(i-1) - pstar(i-1)))
ELSE
epf(i) = p(i-1)
END IF

!Calculate charist price expectation:
epc(i) = p(i-1) + (rho * (epc(i-1) - p(i-2)))
> + ((1 - rho) * (beta * (p(i-1) - p(i-2))))

!Calculate market clearing price:
IF (ABS((muf(i)*varf(i))-(mu*varc(i)))<=EPSILON(0.0D0)) THEN
P(i) = (((wf(i) * epf(i)) + (wc(i) * epc(i))
> - (s(i) * muf(i) * varf(i))) /
> ((1 + r) * (wf(i) + wc(i)))) + eta(i-2)
ELSE
p(i) = (((wf(i) * epf(i) * mu * varc(i))
> + (wc(i) * epc(i) * muf(i) * varf(i))
> - (mu * muf(i) * varf(i) * mu * varc(i) * s(i)))
> / ((1 + r) * ((wf(i) * mu * varc(i))
> + (wc(i) * muf(i) * varf(i)))) + eta(i-2)
END IF

END DO loop1

RETURN
END SUBROUTINE

```

```

SUBROUTINE genpstar (t, seed, pstar, profits)

```

```

!Purpose: end of column 70-->
! To generate bootstapped profits series based on the regression
! dlnrp(t) = 0.006689 - 0.260975dlnrp(t-1) + u(t)

```

```

!
!Record of revisions:
! Date Programmer Description of change
! =====
! 12/06/07 Rhys ap Gwilym Original code
! 14/06/07 Rhys ap Gwilym Convert to subroutine

IMPLICIT NONE

INTEGER, PARAMETER :: double=8 !double precision

!Declare calling parameters:
INTEGER, INTENT(IN) :: t !number of simulation periods
INTEGER, INTENT(IN) :: seed !seed for random number generators
REAL(KIND=double), INTENT(OUT), DIMENSION(t+2) :: pstar, profits !vector of
! fundamental price series

!Declare local variables:
INTEGER, PARAMETER :: it=1000 !number of iterations
INTEGER :: i,j !loop variables
REAL(KIND=double), DIMENSION(174) :: residuals !residuals from actual
! regression
INTEGER, DIMENSION(t+2) :: random !vector of random integers
EXTERNAL G05DYF, G05CBF !random number generators
INTEGER :: G05DYF !random number generators
REAL(KIND=double), DIMENSION(t+2) :: u, dlgrp !vector of random error terms,
! generated dlgrp
REAL(KIND=double), PARAMETER :: c=0.008434957528795 !dlgrp regression
! constant
REAL(KIND=double), PARAMETER :: ar=-0.260974680648267 !dlgrp regression ar1
! parameter
REAL(KIND=double) :: edln, eln !expected dlgrp, lgrp
REAL(KIND=double), PARAMETER :: beta=0.99 !discount factor
REAL(KIND=double) :: discount !cumulative discount factor

DATA residuals /
>-0.028244277,-0.0062064057,-0.011598109,
>0.096912936,-0.048749864,-0.063979229,
>0.021223778,0.04918006,-0.051853865,
>-0.055090369,-0.039011901,0.0074308862,
>0.02776317,-0.064727678,0.040706301,
>0.07420187,-0.10925488,-0.066119839,

```

>0.037778505,0.048828162,0.037724215,
>-0.052218183,0.029925076,-0.0089474608,
>-0.030926285,0.010325857,-0.06479717,
>-0.044803086,-0.018870103,0.12261226,
>0.0036853552,-0.047093562,0.014700671,
>0.080885062,0.0088680662,-0.030635158,
>-0.026921937,0.10712826,0.14938749,
>-0.06694645,-0.11976774,0.058023423,
>-0.20035847,0.01022806,-0.04356883,
>-0.033365071,-0.077630221,-0.11084529,
>-0.075054114,0.12849141,0.041452921,
>-0.06737936,0.0074815804,0.17773763,
>0.1144986,0.036990421,-0.020314605,
>0.020373271,0.031074458,-0.029163859,
>-0.0017408722,-0.032064281,-0.036765652,
>0.10278902,-0.05111146,-0.0030760222,
>-0.025792926,0.0068880323,-0.12411376,
>-0.041866315,-0.064703752,0.011462214,
>0.052449414,0.045696622,-0.077245079,
>0.13868803,0.011173646,0.033479654,
>-0.0067745238,0.068771582,0.039219523,
>0.014729704,-0.046920194,-0.019393053,
>0.044904784,0.059453969,0.041066923,
>-0.025979353,0.00136561,0.012645093,
>-0.00023560257,-0.045974855,-0.027863518,
>0.0044994725,0.064835779,-0.041343831,
>0.046233334,0.019214307,0.027753447,
>-0.052228446,0.040148513,-0.0028085386,
>-0.0038162519,0.016598741,0.032279573,
>-0.055633746,-0.038642745,0.047144406,
>-0.045087208,-0.11398978,-0.05391175,
>-0.020068152,0.031911075,0.0041931567,
>-0.0098681635,-0.01977683,-0.028443641,
>0.069406952,0.0081819931,0.050818787,
>0.0065876972,-0.002110552,0.079210355,
>0.020063703,0.022556511,-0.019174155,
>0.0071703439,-0.04176948,0.046091376,
>-0.015478813,0.029629145,-0.0073499661,
>0.019319379,0.029268053,-0.0011111588,
>-0.024149462,-0.0065146387,-0.0033174068,
>-0.0048189316,-0.019015305,0.051698179,
>-0.11156888,-0.10435068,0.058437885,
>0.11548093,-0.0053956046,-0.050071373,
>-0.040982477,0.015895658,-0.045825347,
>0.015206307,-0.097573309,0.1131244,

```

>-0.095943147,0.055005955,0.014888195,
>0.060069263,0.010329354,0.029832735,
>-0.033858288,0.030936261,-0.044646257,
>0.070313117,0.010553599,0.011350235,
>-0.045893987,-0.020119242,0.0097048307,
>-0.022750675,-0.028034134,0.018347715,
>0.027390806,0.025790253,-0.024368407/

```

```

!generate vector of random integers
CALL GO5CBF(seed)
DO i=1,t+2
random(i) = GO5DYF (1,174)
END DO

```

```

!generate errors randomly from residuals
DO i=1,t+2
u(i) = residuals(random(i))
END DO

```

```

!generate dlnrp series
dlnrp(1) = c + (ar*0) + u(1)
DO i=2,t+2
dlnrp(i) = c + (ar * dlnrp(i-1)) + u(i)
END DO

```

```

!generate profits series
profits(1) = 1.0
DO i=2,t+2
profits(i) = EXP(dlnrp(i) + LOG(profits(i-1)))
END DO

```

```

!generate fundamental price series
DO j=1,t+2
pstar(j) = profits(j)
edln = dlnrp(j)
eln = LOG(profits(j))
DO i=1,it
edln = (edln * ar) + c
pstar(j) = pstar(j) + ((beta**i) * EXP(edln + eln))
eln = eln + edln
END DO
END DO
DO j=1,t+2

```

```
pstar(j) = LOG(pstar(j))
END DO
```

```
RETURN
END SUBROUTINE
```

```
SUBROUTINE moments (t, series, moms)
```

```
!Purpose: end of column 70-->
! To calculate the moments of a series
!
!Record of revisions:
! Date Programmer Description of change
! ==== =====
! 04/07/07 Rhys ap Gwilym Original code
```

```
IMPLICIT NONE
```

```
INTEGER, PARAMETER :: double=8 !double precision
```

```
!Declare calling parameters:
INTEGER, INTENT(IN) :: t !size of series
REAL(KIND=double), INTENT(IN), DIMENSION(t) :: series !seed for random
! number generators
REAL(KIND=double), INTENT(OUT), DIMENSION(4) :: moms !vector of fundamental
! price series
```

```
!Declare local variables:
INTEGER :: i,j !loop variables
```

```
!Calculate mean
moms(1) = 0.0
DO i=1,175
moms(1) = moms(1) + series(i)
END DO
moms(1) = moms(1) / REAL(t)
```

```
!Calculate variance
moms(2) = 0.0
```

```

DO i=1,175
moms(2) = moms(2) + ((series(i) - moms(1))**2)
END DO
moms(2) = moms(2) / (REAL(t)-1)

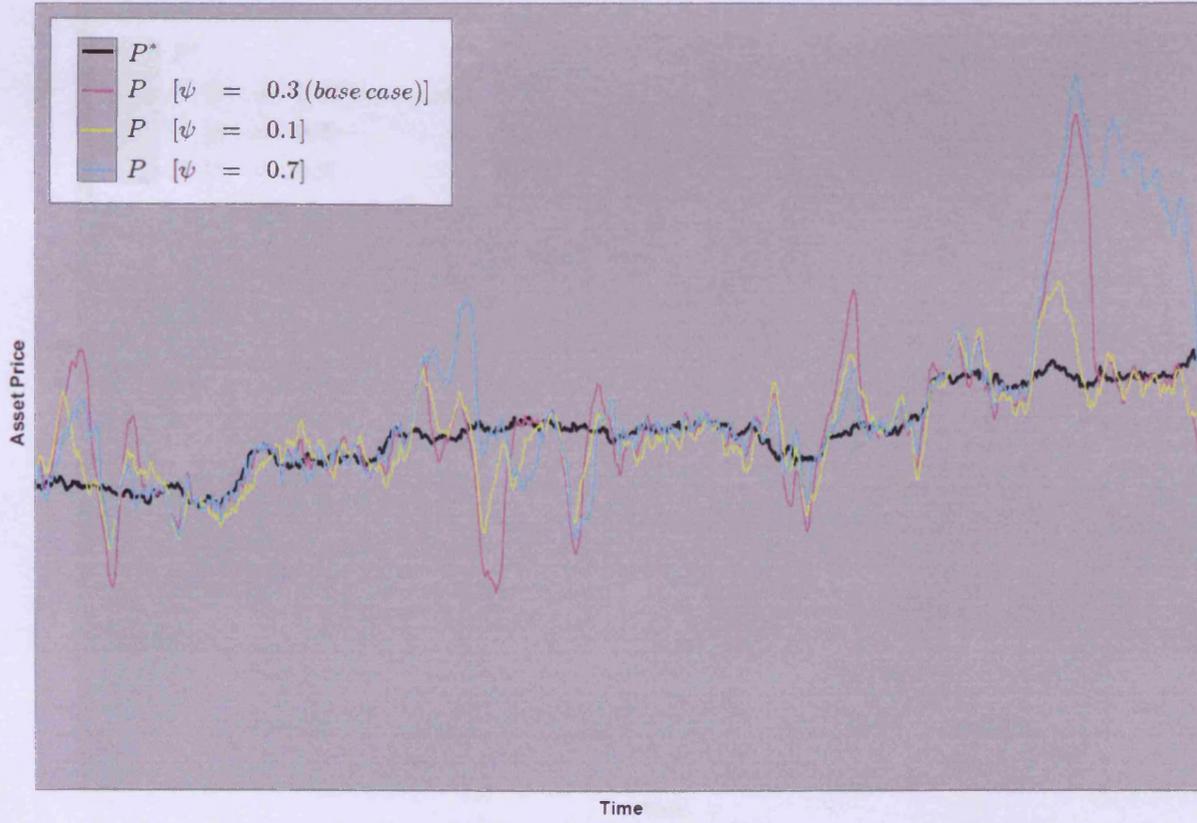
!Calculate skewness
moms(3) = 0.0
DO i=1,175
moms(3) = moms(3) + ((series(i) - moms(1))**3)
END DO
moms(3) = ((REAL(t) / ((REAL(t)-1) * (REAL(t)-2))) * moms(3))
> / (moms(2) ** (1.5))

!Calculate mean
moms(4) = 0.0
DO i=1,175
moms(4) = moms(4) + ((series(i) - moms(1))**4)
END DO
moms(4) = (((((REAL(t)+1) * REAL(t)) / ((REAL(t)-1) * (REAL(t)-2)
> * (REAL(t)-3))) * moms(4)) / (moms(2) ** 2.0))
> - ((3.0 * ((REAL(t)-1) ** 2.0))
> / ((REAL(t)-2) * (REAL(t)-3)))

RETURN
END SUBROUTINE

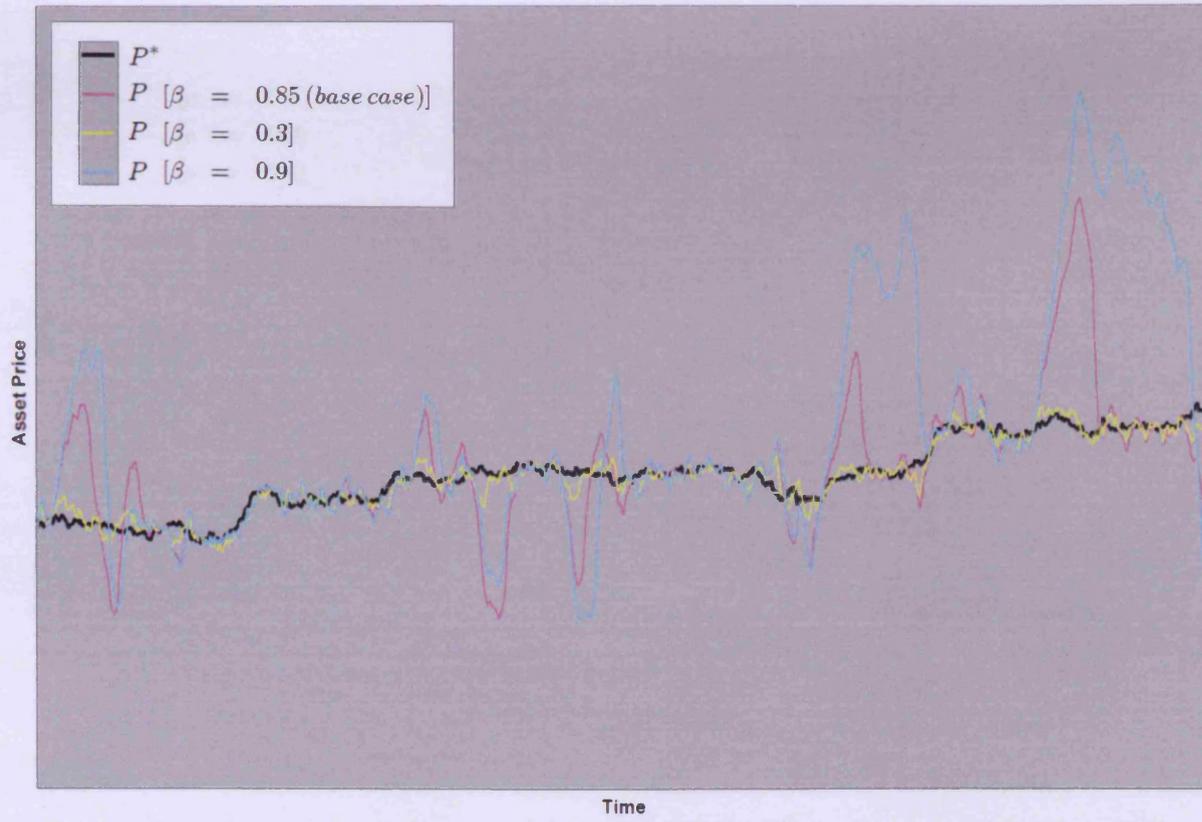
```

Sensitivity to ψ

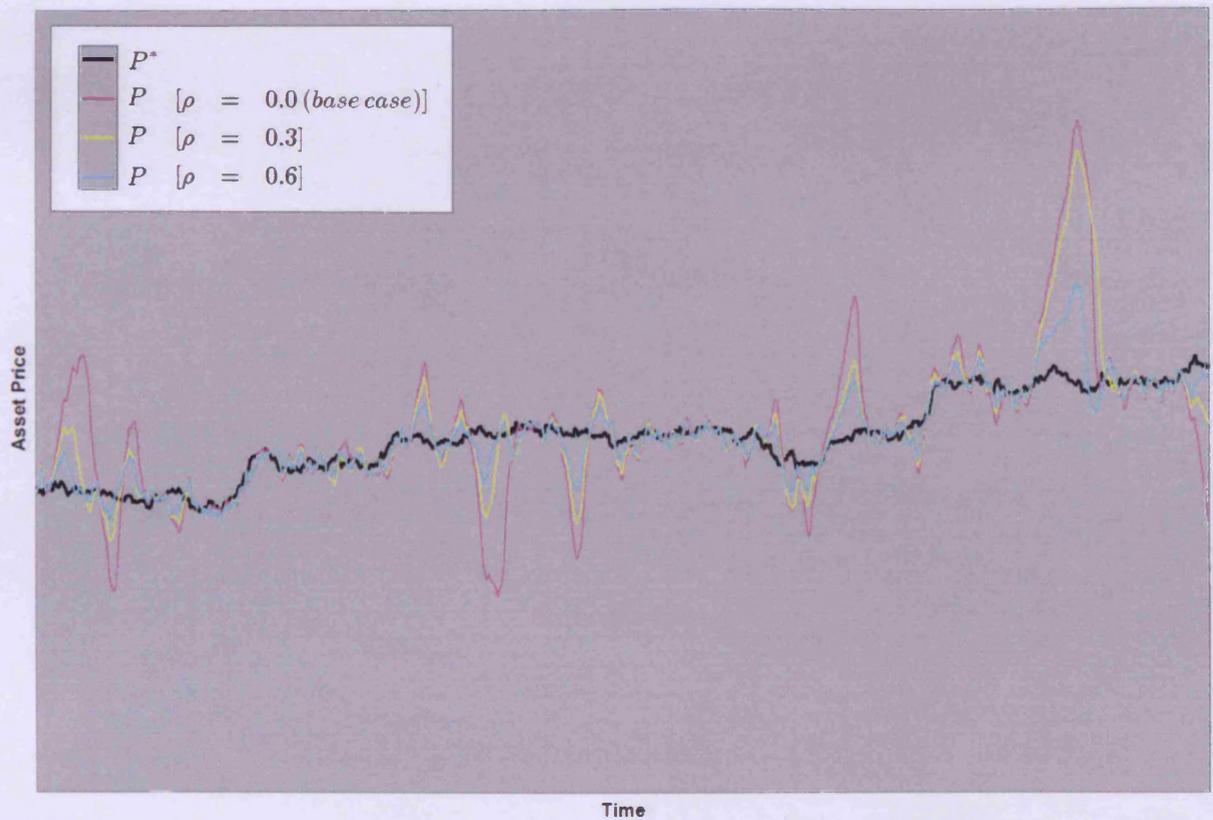


A.2. Stochastic simulations

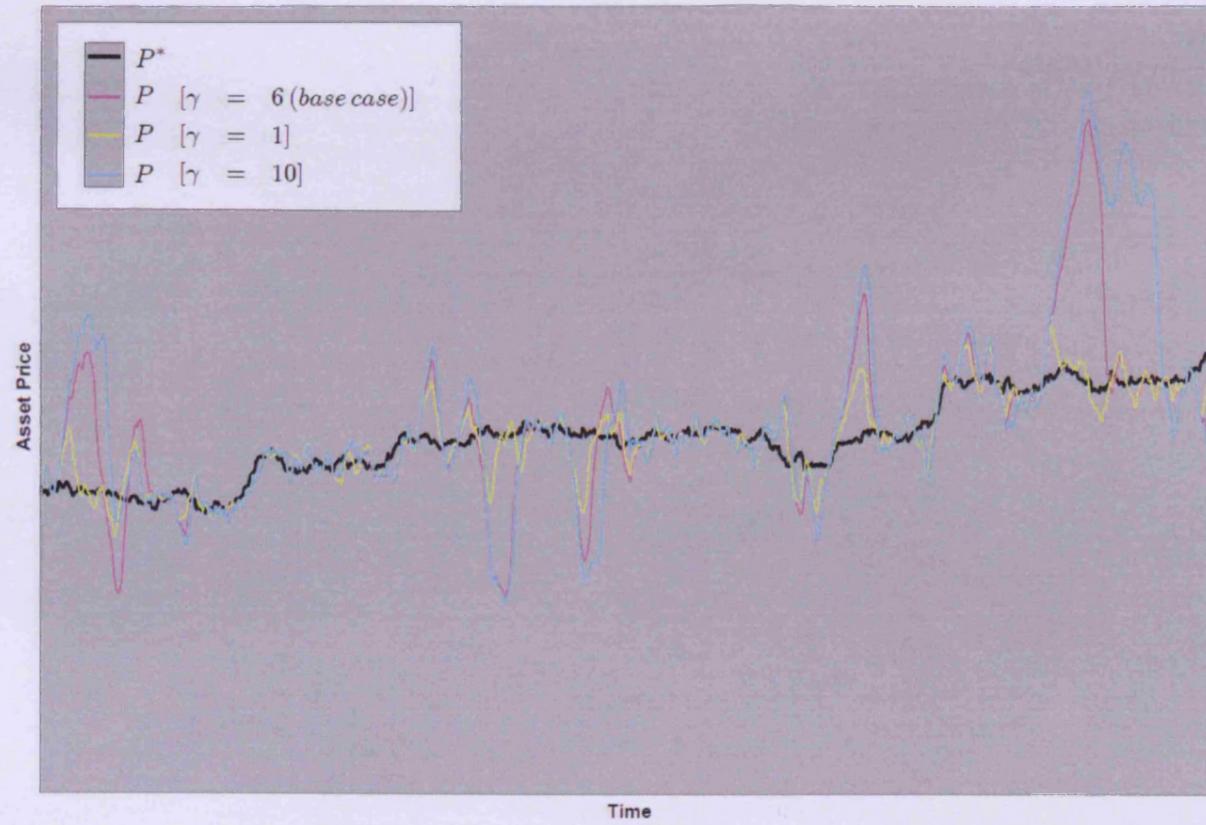
Sensitivity to β



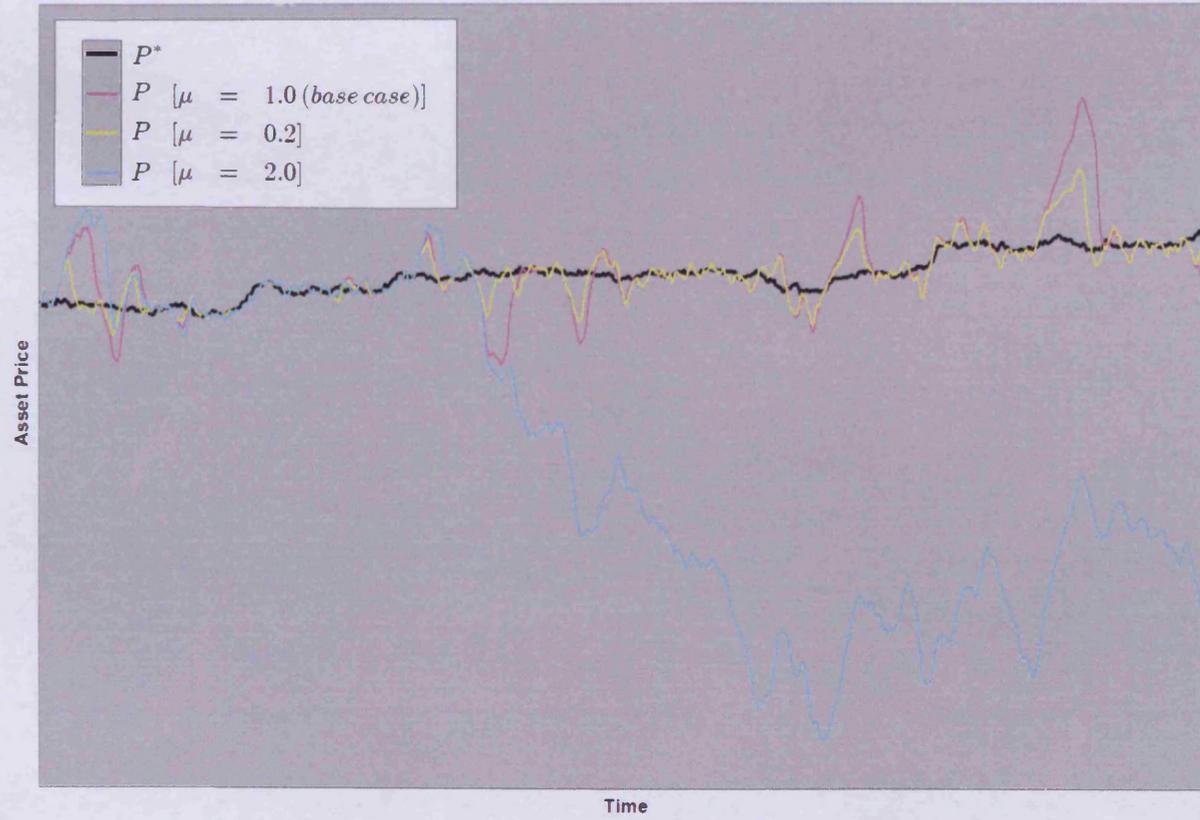
Sensitivity to ρ



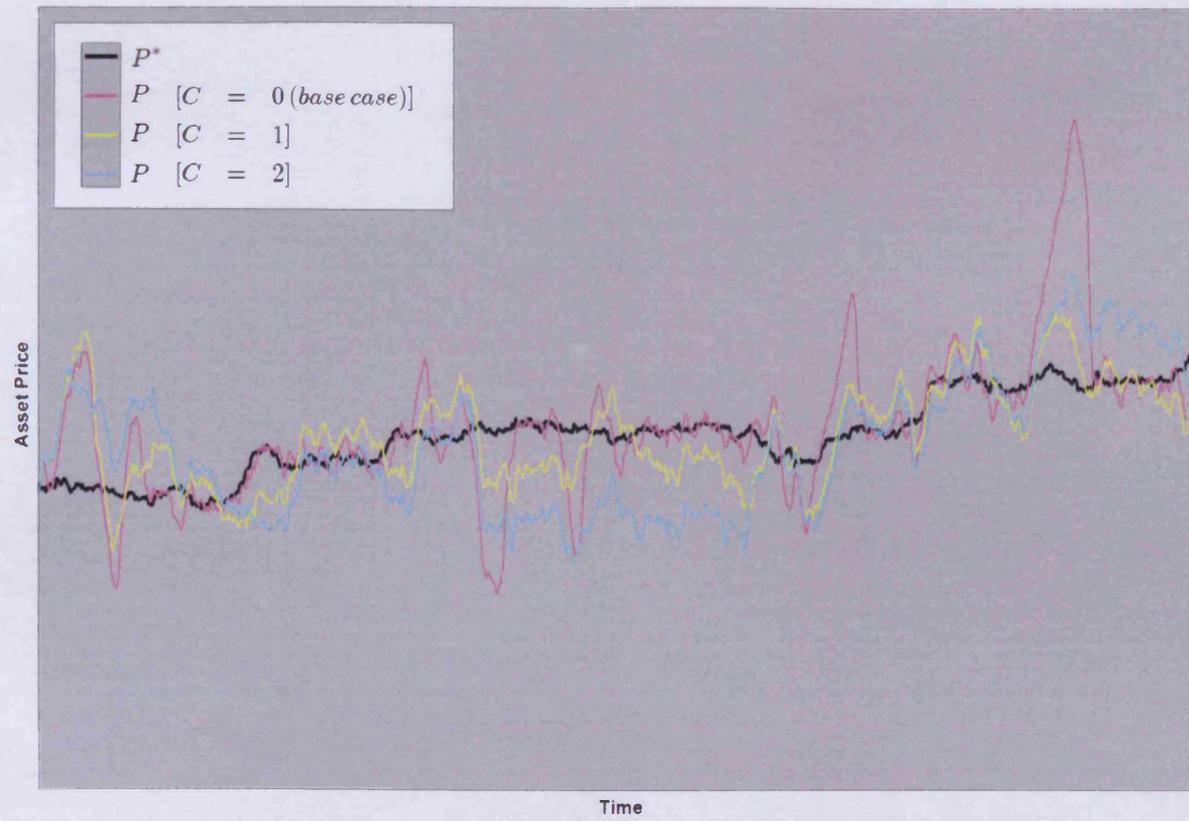
Sensitivity to γ



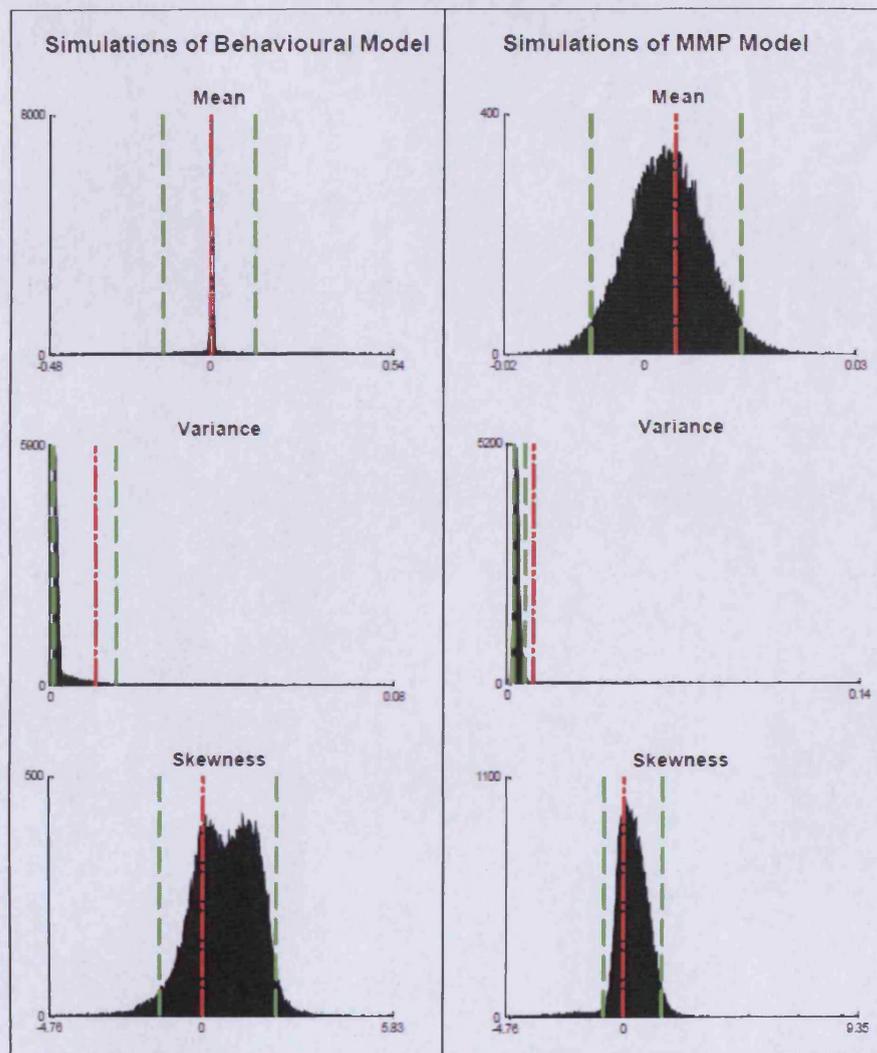
Sensitivity to μ

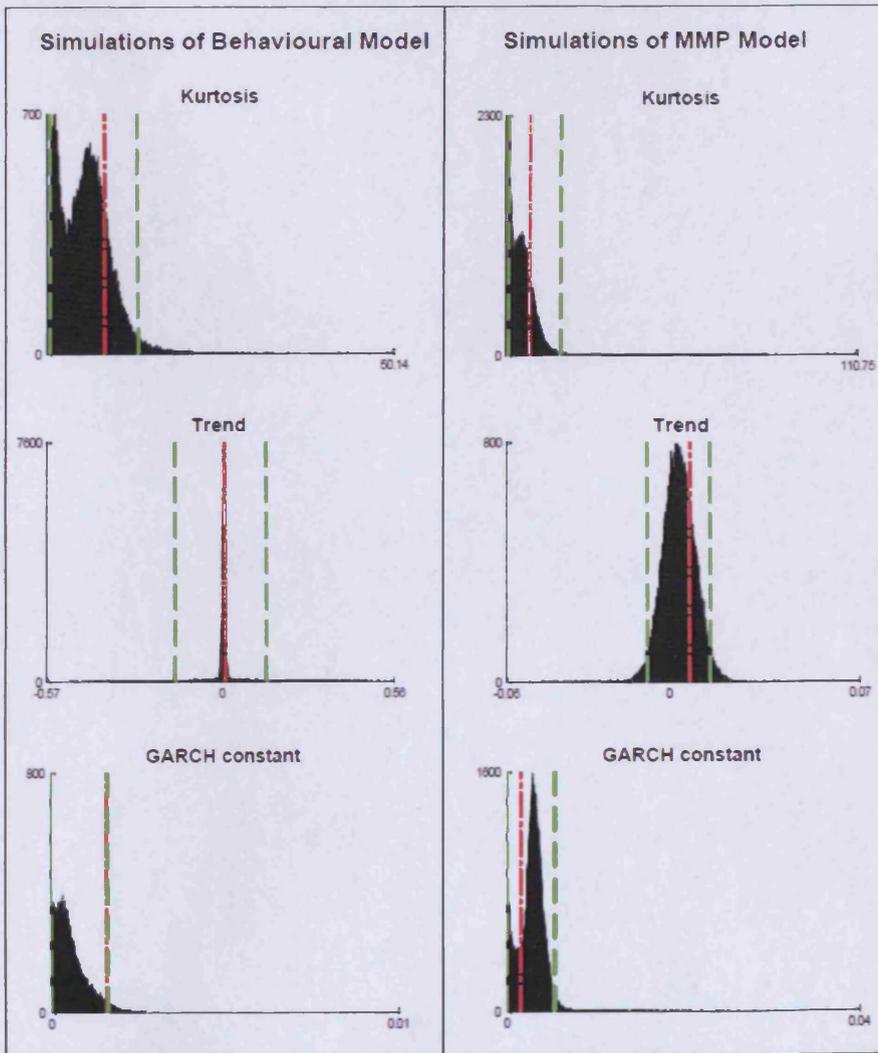


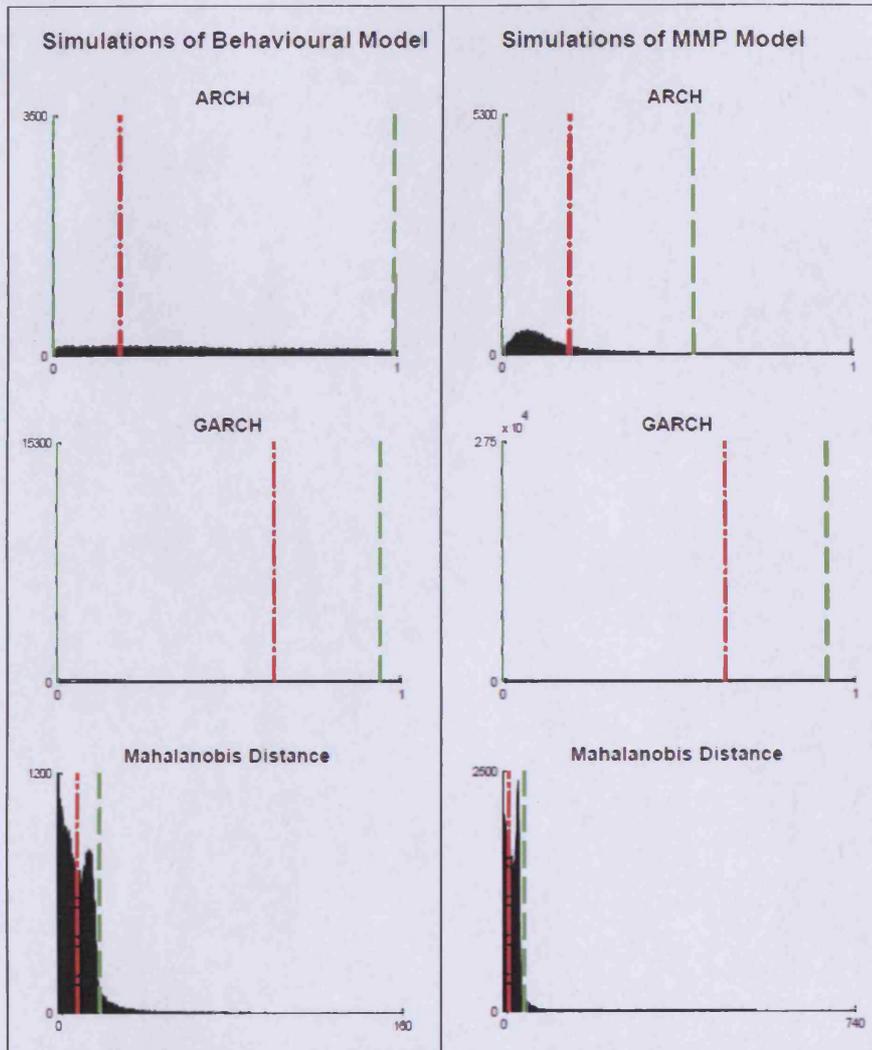
Sensitivity to C



A.3. Results of the Stochastic Simulations







APPENDIX B

A DSGE Model with Asset Bubbles

B.1. Parameters for the Benchmark Model

Intertemporal discount factor:	β	=	0.99
Elasticity in the goods aggregator:	θ	=	7.00
Work disutility coefficient = $1 + \phi$,	ϕ	=	3.00
Probability Calvo fairy does not visit price setter:	η	=	0.67
Probability Calvo fairy does not visit wage setter:	ω	=	0.75
Inertia in productivity shock:	ρ	=	0.95
Elasticity in the labour aggregator:	γ	=	7.00
Depreciation rate:	δ	=	0.025
Coefficient in adjustment cost for investment:	ν	=	8.00
Share of capital in the production function:	α	=	0.25
Chartist rule parameter:	χ_c	=	$\frac{0.99}{1.99}$
Fundamentalist rule parameter:	χ_f	=	0.50
Uncertainty bound in fundamentalist rule:	C	=	0.00
Propensity to switch between forecasting rules:	v	=	3.75
Taylor rule weight on inflation:	ζ_Π	=	2.02
Taylor rule weight on output gap:	ζ_Y	=	0.184
Taylor rule weight on asset mis-pricing:	ζ_Q	=	various
Standard deviation of productivity shock:	σ_ϵ	=	0.0086

B.2. The Dynare Sub-model: Equation listing

B.2.1. Rational variables

$$\begin{aligned}
\frac{1}{c_t^R} &= \lambda_t^R \\
w_t^{*R^{1+\gamma\phi}} &= \frac{\gamma}{\gamma-1} \frac{wb_t^R}{wa_t^R} \\
wb_t^R &= l_t^{1+\phi} + \omega\beta E_t \left[\left(\frac{1 + \Pi_{t+1}^R w_{t+1}^R}{1 + \Pi} \frac{w_t^R}{w_t^R} \right)^{\gamma(1+\phi)} wb_{t+1}^R \right] \\
wa_t^R &= \lambda_t^R w_t^R l_t^R + \omega\beta E_t \left[\left(\frac{1 + \Pi_{t+1}^R w_{t+1}^R}{1 + \Pi} \frac{w_t^R}{w_t^R} \right)^{\gamma-1} wa_{t+1}^R \right] \\
1 &= (1-\omega)w_t^{*R^{1-\gamma}} + \omega \left(\frac{1 + \Pi_t^R w_t^R}{1 + \Pi} \frac{w_t^R}{w_{t-1}^R} \right)^{\gamma-1} \\
\mu_t^R &= \beta E_t \left\{ \lambda_{t+1}^R r_{t+1}^{k^*R} + \mu_{t+1}^R \left[1 - \delta + \frac{\nu}{2} \left[\left(\frac{i_{t+1}^R}{k_{t+1}^R} \right)^2 - \delta^2 \right] \right] \right\} \\
\lambda_t^R &= \mu_t^R \left\{ 1 - \nu \left[\frac{i_t^R}{k_t^R} - \delta \right] \right\} \\
E_t r_{t+1}^{bR} &= \frac{\lambda_t^R}{\beta E_t \lambda_{t+1}^R} - 1 \\
q_t^{*R} &= E_t \beta \frac{\lambda_{t+1}^R}{\lambda_t^R} [r_{t+1}^{kR} + q_{t+1}^{*R}] \\
k_{t+1}^R &= (1-\delta)k_t^R + i_t^R - \frac{\nu}{2} \left(\frac{i_t^R}{k_t^R} - \delta \right)^2 k_t^R \\
\frac{k_t^R}{l_t^R} &= \frac{\alpha}{1-\alpha} \frac{w_t^R}{r_t^{kR}} \\
mc_t^R &= \frac{w_t^{R^{1-\alpha}} r_t^{kR\alpha}}{z_t \alpha^\alpha (1-\alpha)^{1-\alpha}}
\end{aligned}$$

$$\begin{aligned}
p_t^{*R} &= \frac{\theta}{\theta - 1} \frac{pb_t^R}{pa_t^R} \\
pb_t^R &= \lambda_t^R mc_t^R y_t^R + \eta\beta \left(\frac{1 + E_t \Pi_{t+1}^R}{1 + \Pi} \right)^\theta E_t pb_{t+1}^R \\
pa_t^R &= \lambda_t^R y_t^R + \eta\beta \left(\frac{1 + E_t \Pi_{t+1}^R}{1 + \Pi} \right)^{\theta-1} E_t pa_{t+1}^R \\
1 &= (1 - \eta) p_t^{*R^{1-\theta}} + \eta \left(\frac{1 + \Pi}{1 + \Pi_t^R} \right)^{1-\theta} \\
y_t^R &= \frac{z_t k_t^{R\alpha} l_t^{R(1-\alpha)}}{pd_t^R} \\
pd_t^R &= (1 - \eta) p_t^{*R^{-\theta}} + \eta \left(\frac{1 + \Pi_t^R}{1 + \Pi} \right)^\theta pd_{t-1}^R \\
y_t^R &= c_t^R + i_t^R \\
1 + r_t^{bR} &= \frac{1 + R_t^{bR}}{1 + \Pi_t^R} \\
R_{t+1}^{bR} &= r^{b*} + \Pi_t^R + \zeta_\Pi (\Pi_t^R - \Pi^*) + \zeta_y \ln \left(\frac{y_t^R}{y^*} \right) + \varepsilon_t^{MP} \\
wd_t^R &= (1 - \omega) w_t^{*R^{-\gamma(1+\phi)}} + \omega \left(\frac{w_t^R}{w_{t-1}^R} \frac{1 + \Pi_t^R}{1 + \Pi} \right)^{\gamma(1+\phi)} wd_{t-1}^R \\
U_t^R &= \ln(c_t^R) - \frac{l_t^{R(1+\phi)} wd_t^R}{1 + \phi} + \beta E_t U_{t+1}^R
\end{aligned}$$

B.2.2. Behavioural variables:

$$\begin{aligned}
\frac{1}{c_t^B} &= \lambda_t^B \\
w_t^{*B^{1+\gamma\phi}} &= \frac{\gamma}{\gamma-1} \frac{wb_t^B}{wa_t^B} \\
wb_t^B &= l_t^{1+\phi} + \omega\beta E_t \left[\left(\frac{1 + \Pi_{t+1}^B w_{t+1}^B}{1 + \Pi} \frac{w_{t+1}^B}{w_t^B} \right)^{\gamma(1+\phi)} wb_{t+1}^B \right] \\
wa_t^B &= \lambda_t^B w_t^B l_t^B + \omega\beta E_t \left[\left(\frac{1 + \Pi_{t+1}^B w_{t+1}^B}{1 + \Pi} \frac{w_{t+1}^B}{w_t^B} \right)^{\gamma-1} wa_{t+1}^B \right] \\
1 &= (1-\omega)w_t^{*B^{1-\gamma}} + \omega \left(\frac{1 + \Pi_t^B w_t^B}{1 + \Pi} \frac{w_t^B}{w_{t-1}^B} \right)^{\gamma-1} \\
\mu_t^B &= \beta E_t \left\{ \lambda_{t+1}^B r_{t+1}^{k^*B} + \mu_{t+1}^B \left[1 - \delta + \frac{\nu}{2} \left[\left(\frac{i_{t+1}^B}{k_{t+1}^B} \right)^2 - \delta^2 \right] \right] \right\} \\
\lambda_t^B &= \mu_t^B \left\{ 1 - \nu \left[\frac{i_t^B}{k_t^B} - \delta \right] \right\} \\
E_t r_{t+1}^{bB} &= \frac{\lambda_t^B}{\beta E_t \lambda_{t+1}^B} - 1 \\
q_t^B &= E_t \beta \frac{\lambda_{t+1}^B}{\lambda_t^B} [r_{t+1}^{kB} + q_t^{BF}] \\
q_t^B &= q_t^{*R} + \varepsilon_t^q \\
q_t^{BF} &= E_t q_{t+1}^{*R} + \varepsilon_t^{qF} \\
k_{t+1}^B &= (1-\delta)k_t^B + i_t^B - \frac{\nu}{2} \left(\frac{i_t^B}{k_t^B} - \delta \right)^2 k_t^B \\
\frac{k_t^B}{l_t^B} &= \frac{\alpha}{1-\alpha} \frac{w_t^B}{r_t^{kB}} \\
mc_t^B &= \frac{w_t^{B^{1-\alpha}} r_t^{kB\alpha}}{z_t \alpha^\alpha (1-\alpha)^{1-\alpha}}
\end{aligned}$$

$$\begin{aligned}
p_t^{*B} &= \frac{\theta}{\theta - 1} \frac{pb_t^B}{pa_t^B} \\
pb_t^B &= \lambda_t^B m c_t^B y_t^B + \eta \beta \left(\frac{1 + E_t \Pi_{t+1}^B}{1 + \Pi} \right)^\theta E_t pb_{t+1}^B \\
pa_t^B &= \lambda_t^B y_t^B + \eta \beta \left(\frac{1 + E_t \Pi_{t+1}^B}{1 + \Pi} \right)^{\theta-1} E_t pa_{t+1}^B \\
1 &= (1 - \eta) p_t^{*B^{1-\theta}} + \eta \left(\frac{1 + \Pi}{1 + \Pi_t^B} \right)^{1-\theta} \\
y_t^B &= \frac{z_t k_t^{B\alpha} l_t^{B(1-\alpha)}}{pd_t^B} \\
pd_t^B &= (1 - \eta) p_t^{*B-\theta} + \eta \left(\frac{1 + \Pi_t^B}{1 + \Pi} \right)^\theta pd_{t-1}^B \\
y_t^B &= c_t^B + i_t^B \\
1 + r_t^{bB} &= \frac{1 + R_t^{bB}}{1 + \Pi_t^B} \\
R_{t+1}^{bB} &= r^{b*} + \Pi_t^B + \zeta_\Pi (\Pi_t^B - \Pi^*) + \zeta_y \ln \left(\frac{y_t^B}{y^*} \right) + \varepsilon_t^{MP} \\
wd_t^B &= (1 - \omega) w_t^{*B-\gamma(1+\phi)} + \omega \left(\frac{w_t^B}{w_{t-1}^B} \frac{1 + \Pi_t^B}{1 + \Pi} \right)^{\gamma(1+\phi)} wd_{t-1}^B \\
U_t^B &= \ln(c_t^B) - \frac{l_t^{B(1+\phi)} wd_t^B}{1 + \phi} + \beta E_t U_{t+1}^B
\end{aligned}$$

B.2.3. Productivity process:

$$\ln(z_t) = \varrho \ln(z_{t-1}) + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma_\varepsilon)$$

B.3. The Dynare Code

```

// DYNARE MODEL
// NOTATION:
// The following are in log form:
// z, zq, zfq
// r_c, r_lambda, r_wstar, r_l, r_w, r_mu, r_i, r_k, r_q, r_mc, r_pstar, r_y
// b_c, b_lambda, b_wstar, b_l, b_w, b_mu, b_i, b_k, b_q, b_mc, b_pstar, b_y
// rf_q, bf_q
// The following are log of the gross rate:
// r_rb, r_RBi, r_pi
// b_rb, b_RBi, b_pi
// The following are in levels:
// r_rk, r_wb, r_wa, r_pb, r_pa, r_pd, r_wd, r_v, r_u, r_wdg
// b_rk, b_wb, b_wa, b_pb, b_pa, b_pd, b_wd, b_v, b_u, b_wdg
// rf_rk, bf_rk
//VARIABLES
var z, zq, zfq,
    r_c, r_lambda, r_wstar, r_l, r_w, r_mu, r_i, r_k, r_q, r_mc, r_pstar, r_y,
    b_c, b_lambda, b_wstar, b_l, b_w, b_mu, b_i, b_k, b_q, b_mc, b_pstar, b_y,
    rf_q, bf_q, rf_lambda, bf_lambda,
    r_rb, r_RBi, r_pi,
    b_rb, b_RBi, b_pi,
    r_rk, r_wb, r_wa, r_pb, r_pa, r_pd, r_wd, r_v, r_u,
    b_rk, b_wb, b_wa, b_pb, b_pa, b_pd, b_wd, b_v, b_u,
    rf_rk, bf_rk;
varexo eps, eq, efq, pol;
//PARAMETERS
parameters gamma, phi, omega, beta, sspi, delta, nu, alpha, theta, eta, rho,
    RBstar, zetapi, pistar, zetay, ystar, rhoq, rhofq, sigma;
gamma = 7.0;
phi = 3.0;
omega = 0.67;
beta = 0.99;
sspi = 0.0198;
delta = 0.025;
nu = 8.0;
alpha = 0.25;
theta = 7.0;
eta = 0.67;
rho = 0.95;
zetapi = 0.5;
pistar = 0.0198;
zetay = 0.5;

```

```

ystar = 0.495432665018909;
rhoq = 0.0;
rhofq = 0.0;
sigma = 0.0001;
//MODEL EQUATIONS: [NB k=k(t+1), rb=rb(t+1), RBi=Rb(t+1), rkstar=rkstar(t+1)]
model;
//RATIONAL MODEL
//Household optimisation:
//marginal utility of consumption set equal to marginal utility of wealth:
r_c = -r_lambda;
//Calvo wage setting. Expected wage equals markup over expected disutility
of work:
r_wstar*(1+gamma*phi) = log(gamma)-log(gamma-1)+log(r_wb)-log(r_wa);
r_wb = exp(r_l*(1+phi))+omega*beta*r_wb(+1)*(exp(r_pi(+1)-sspi+r_w(+1)-r_w))
^(gamma*(1+phi));
r_wa = exp(r_lambda+r_w+r_l)+omega*beta*r_wa(+1)*(exp(r_pi(+1)-sspi+r_w(+1)-
r_w))^(gamma-1);
1 = (1-omega)*exp(r_wstar*(1-gamma))+omega*(exp(r_pi-sspi+r_w-r_w(-1)))^
(gamma-1);
//expected returns to holding capital and bonds equalised, but with a
behavioural wedge:
exp(r_mu) = beta*(exp(r_lambda(+1))*r_rk(+1)+exp(r_mu(+1))*(1-delta+(nu/2)*
(exp(2*(r_i(+1)-r_k))-delta^2)));
r_lambda = r_mu+log(1-nu*(exp(r_i-r_k(-1))-delta));
r_rb = r_lambda-r_lambda(+1)-log(beta);
// r_rkstar/r_rk(+1) = 1+wdg;
// wdg = wdg+erk;
exp(r_q) = beta*(exp(r_lambda(+1)-r_lambda)*r_rk(+1)+exp(r_lambda(+1)-
r_lambda+r_q(+1)));
//Law of motion of capital:
exp(r_k) = (1-delta)*exp(r_k(-1))+exp(r_i)-(nu/2)*((exp(r_i-r_k(-1))-delta)
^2)*exp(r_k(-1));
//Firm optimisation:
//Factor input ratio set equal to ratio of marginal returns:
r_k(-1)-r_l = log(alpha)-log(1-alpha)+r_w-log(r_rk);
//Calvo pricing. xpctd pric quals markup ovr xpctd marginal cost:
r_mc = r_w*(1-alpha)+log(r_rk)*alpha-z-alpha*log(alpha)-(1-alpha)*log(1-
alpha);
exp(r_pstar) = (theta/(theta-1))*(r_pb/r_pa);
r_pb = exp(r_lambda+r_mc+r_y)+eta*beta*((exp(r_pi(+1)-sspi))^theta)*
r_pb(+1);
r_pa = exp(r_lambda+r_y)+eta*beta*((exp(r_pi(+1)-sspi))^(theta-1))*
r_pa(+1);
1 = (1-eta)*(exp(r_pstar)^(1-theta))+eta*((exp(sspi-r_pi))^(1-theta));
//Production and market clearing:

```

```

r_y = z+(r_k(-1)*alpha)+(r_l*(1-alpha))-log(r_pd);
r_pd = (1-eta)/exp(r_pstar*theta)+eta*((exp(r_pi-sspi))^theta)*r_pd(-1);
z = rho*z(-1)+eps;
exp(r_y) = exp(r_c)+exp(r_i);
//Definition of real interest rate and monetary policy:
r_rb = r_RB_i-r_pi;
// RB_i = pistar-log(beta)+zetapi*(pi-pistar)+zetay*(y-ystar)+pol;
// RB_i = pi(-1)-log(beta);//+zetapi*(pi(-1)-pistar)+zetay*(y-ystar)+pol;
// r_pi = 0.0198;
//Canzoneri:
// RB_i = -(1-0.824)*log(beta)+0.824*RB_i(-1)+(1-0.824)*2.02*pi+(1-0.824)*
    0.184*(y-ystar)+pol;
//Canzoneri no smoothing:
r_RB_i = r_pi-log(beta)+2.02*(r_pi-pistar)+0.184*(r_y-ystar)+pol;
//Trad TR:
// RB_i = -log(beta)+pi+0.5*(pi-sspi)+0.5*(y-ystar)+pol;
//Welfare:
r_wd = (1-omega)/(exp(r_wstar*gamma*(1+phi)))+omega*((exp(r_w-r_w(-1))*(exp
    (r_pi-sspi))^gamma*(1+phi)))*r_wd(-1);
r_v = r_c-exp(r_l*(1+phi))*r_wd/(1+phi);
r_u = r_v+beta*r_u(+1);
//Further variables for export:
rf_lambda = r_lambda(+1);
rf_q = r_q(+1);
rf_rk = r_rk(+1);
//BEHAVIOURAL MODEL
//Household optimisation:
//marginal utility of consumption set equal to marginal utility of wealth:
b_c = -b_lambda;
//Calvo wage setting. Expected wage equals markup over expected disutility
    of work:
b_wstar*(1+gamma*phi) = log(gamma)-log(gamma-1)+log(b_wb)-log(b_wa);
b_wb = exp(b_l*(1+phi))+omega*beta*b_wb(+1)*(exp(b_pi(+1)-sspi+b_w(+1)-b_w))
    ^gamma*(1+phi);
b_wa = exp(b_lambda+b_w+b_l)+omega*beta*b_wa(+1)*(exp(b_pi(+1)-sspi+b_w(+1)-
    b_w))^gamma*(1+phi);
1 = (1-omega)*exp(b_wstar*(1-gamma))+omega*(exp(b_pi-sspi+b_w-b_w(-1)))^
    (gamma-1);
//expected returns to holding capital and bonds equalised, but with a
    behavioural wedge:
exp(b_mu) = beta*(exp(b_lambda(+1))*bf_rk+exp(b_mu(+1))*(1-delta+(nu/2)*(exp
    (2*(b_i(+1)-b_k))-delta^2)));
b_lambda = b_mu+log(1-nu*(exp(b_i-b_k(-1))-delta));
b_rb = b_lambda-b_lambda(+1)-log(beta);
// b_rkstar/b_rk(+1) = 1+wdg;

```

```

// wdg = wdg*wdg(-1)+erk;
exp(b_q) = beta*(exp(b_lambda(+1)-b_lambda)*b_rk(+1)+exp(b_lambda(+1)-
      b_lambda+bf_q));
b_q = r_q+zq;
zq = rhoq*zq(-1)+eq;
bf_q = r_q(+1)+zfq;
zfq = rhofq*zfq(-1)+efq;
// bf_q = r_q(+1)+eqf;
//Law of motion of capital:
exp(b_k) = (1-delta)*exp(b_k(-1))+exp(b_i)-(nu/2)*((exp(b_i-b_k(-1))-delta)^
      2)*exp(b_k(-1));
//Firm optimisation:
//Factor input ratio set equal to ratio of marginal returns:
b_k(-1)-b_l = log(alpha)-log(1-alpha)+b_w-log(b_rk);
//Calvo pricing. Expected price equals markup over expected marginal cost:
b_mc = b_w*(1-alpha)+log(b_rk)*alpha-z-alpha*log(alpha)-(1-alpha)*log(1-
      alpha);
exp(b_pstar) = (theta/(theta-1))*(b_pb/b_pa);
b_pb = exp(b_lambda+b_mc+b_y)+eta*beta*((exp(b_pi(+1)-sspi))^theta)*b_pb(+1);
b_pa = exp(b_lambda+b_y)+eta*beta*((exp(b_pi(+1)-sspi))^(theta-1))*b_pa(+1);
1 = (1-eta)*(exp(b_pstar)^(1-theta))+eta*((exp(sspi-b_pi))^(1-theta));
//Production and market clearing:
b_y = z+(b_k(-1)*alpha)+(b_l*(1-alpha))-log(b_pd);
b_pd = (1-eta)/exp(b_pstar*theta)+eta*((exp(b_pi-sspi))^theta)*b_pd(-1);
exp(b_y) = exp(b_c)+exp(b_i);
//Definition of real interest rate and monetary policy:
b_rb = b_RB_i-b_pi;
// RB_i = pi_star-log(beta)+zeta*pi*(pi-pi_star)+zeta*y*(y-y_star)+pol;
// RB_i = pi(-1)-log(beta);//+zeta*pi*(pi(-1)-pi_star)+zeta*y*(y-y_star)+pol;
// b_pi = 0.0198;
//Canzoneri:
// RB_i = -(1-0.824)*log(beta)+0.824*RB_i(-1)+(1-0.824)*2.02*pi+(1-0.824)*0.184
      *(y-y_star)+pol;
//Canzoneri no smoothing:
b_RB_i = b_pi-log(beta)+2.02*(b_pi-pi_star)+0.184*(b_y-y_star)+pol;
//Trad TR:
// RB_i = -log(beta)+pi+0.5*(pi-sspi)+0.5*(y-y_star)+pol;
//Welfare:
b_wd = (1-omega)/(exp(b_wstar*gamma*(1+phi)))+omega*((exp(b_w-b_w(-1))*(exp
      (b_pi-sspi)))^(gamma*(1+phi)))*b_wd(-1);
b_v = b_c-exp(b_l*(1+phi))*b_wd/(1+phi);
b_u = b_v+beta*b_u(+1);
//Further variables for export:
bf_lambda = b_lambda(+1);
end;

```

```

initval;
z = 0.00;
zq = 0.00;
zfq = 0.00;
eps = 0.00;
eq = 0.00;
efq = 0.00;
pol = 0.00;
r_rk = ((1.0-beta)/beta)+delta;
r_rb = -log(beta);
r_pi = pistar;
r_pstar = 0.0;
r_RB_i = r_pi-log(beta);
r_mc = r_pstar+log(theta-1.0)-log(theta);
r_w = (r_mc+z+alpha*log(alpha)+(1.0-alpha)*log(1.0-alpha)-(log(r_rk)*alpha))
      *(1.0/(1.0-alpha)));
r_wstar = 0.0;
r_k = (log(gamma-1.0)-log(gamma)+r_w-phi*(log(1.0-alpha)-log(alpha)+log
      (r_rk)-r_w)-log(exp(z)*(((1.0-alpha)/alpha)*(r_rk/exp(r_w)))^(1.0-
      alpha)-delta))/(1+phi);
r_l = log(1.0-alpha)-log(alpha)+log(r_rk)+r_k-r_w;
r_c = log(exp(z)*(((1.0-alpha)/alpha)*(r_rk/exp(r_w)))^(1.0-alpha)-delta)+
      r_k;
r_y = log(exp(r_c)+(delta*exp(r_k)));
r_i = log(delta)+r_k;
r_lambda = -r_c;
r_pb = exp(r_lambda+r_mc+r_y)/(1-(eta*beta));
r_pa = exp(r_lambda+r_y)/(1-(eta*beta));
r_wb = (exp(r_l)^(1+phi))/(1-(omega*beta));
r_wa = exp(r_lambda+r_w+r_l)/(1-(omega*beta));
r_pd = 1.00000;
r_mu = r_lambda;
r_wd = 1.00000;
r_v = r_c-exp(r_l*(1+phi))/(1+phi);
r_u = (1/(1-beta))*r_v;
r_q = log(beta)+log(r_rk)-log(1-beta);
b_rk = r_rk;
b_rb = r_rb;
b_pi = r_pi;
b_pstar = r_pstar;
b_RB_i = r_RB_i;
b_mc = r_mc;
b_w = r_w;
b_wstar = r_wstar;
b_k = r_k;

```

```
b_l = r_l;
b_c = r_c;
b_y = r_y;
b_i = r_i;
b_lambda = r_lambda;
b_pb = r_pb;
b_pa = r_pa;
b_wb = r_wb;
b_wa = r_wa;
b_pd = r_pd;
b_mu = r_mu;
b_wd = r_wd;
b_v = r_v;
b_u = r_u;
b_q = r_q;
end;
steady;
check;
shocks;
var eps = sigma^2;
var eq = (1*sigma)^2;
var efq = (1*sigma)^2;
var pol = sigma^2;
end;
stoch_simul(periods=2100, irf=40) r_c b_c r_y b_y r_k b_k r_l b_l r_rk b_rk
rf_rk bf_rk r_q b_q rf_q bf_q;
```

B.4. The Behavioural Sub-model: Equation listing

$$\begin{aligned}
E_{c,t}(q_t) &= (1 + \chi_c) q_{t-1} + \chi_c E_{c,t-1}(q_{t-1}) - 2\chi_c q_{t-2} \\
E_{c,t}(q_{t+1}) &= (1 + \chi_c) E_{c,t}(q_t) - \chi_c (1 + \chi_c) E_{c,t-1}(q_{t-1}) + \chi_c E_{c,t-1}(q_t) - \chi_c (1 - \chi_c) q_{t-1} \\
E_{f,t}(q_t) &= (1 - \chi_f) q_{t-1} + \chi_f q_t^* \quad \text{where } |q_{t-1} - q_t^*| > C \\
&= q_{t-1} \quad \text{where } |q_{t-1} - q_t^*| \leq C \\
E_{f,t}(q_{t+1}) &= (1 - \chi_f) E_{f,t}(q_t) + \chi_f q_t^* \quad \text{where } |E_{f,t}(q_t) - q_t^*| > C \\
&= E_{f,t}(q_t) \quad \text{where } |E_{f,t}(q_t) - q_t^*| \leq C \\
\Omega_{c,t} &= [q_{t-1} - q_{t-2} (1 + r_{t-1}^{bB})] \cdot \text{sign} [E_{c,t-2}(q_{t-2}) - q_{t-2}] \\
\Omega_{f,t} &= [q_{t-1} - q_{t-2} (1 + r_{t-1}^{bB})] \cdot \text{sign} [E_{f,t-2}(q_{t-2}) - q_{t-2}] \\
w_{c,t} &= \frac{\exp(\nu \Omega_{c,t})}{\exp(\nu \Omega_{f,t}) + \exp(\nu \Omega_{c,t})} \\
w_{f,t} &= 1 - w_{c,t} \\
q_t &= w_{c,t} E_{c,t}(q_t) + w_{f,t} E_{f,t}(q_t) \\
E_{b,t} q_{t+1} &= w_{c,t} E_{c,t}(q_{t+1}) + w_{f,t} E_{f,t}(q_{t+1}) \\
E_{b,t} r_{t+1}^k &= \frac{q_t}{\beta} \frac{\lambda_t}{E_t \lambda_{t+1}} - E_{b,t} q_{t+1} \\
\varepsilon_t^q &= \ln \left(\frac{q_t}{q_t^*} \right) \\
\varepsilon_t^{qF} &= \ln \left(\frac{E_{b,t} q_{t+1}}{E_t q_{t+1}^*} \right) \\
\varepsilon_t^{MP} &= \zeta_Q \ln \left(\frac{q_{t-1}}{q_{t-1}^*} \right)
\end{aligned}$$

B.5. The MATLAB code

```

%Run dynare model
%dynare d;

t = 40;                                %time periods to calculate
for i=1:endo_nbr;
    varlist(i,:)=lgy_(dr_.order_var(i,:);    %reordered variable list
end;

%Save steady state values to first column of X matrix (re-ordered)
for i=1:endo_nbr
    X(i,1) = ys_(dr_.order_var(i));
end

%Second order approximation is given by:
%y(t) = ys + A yh(t-1) + B u(t) + 0.5C(yh(t-1)(kron*)yh(t-1)) +
%      0.5D(u(t)(kron*)u(t)) + E(yh(t-1)(kron*)u(t))
%where ys is the steady state value of y and yh(t)=y(t)-ys.
%A, B, C, D, E from dynare:
A = [zeros(endo_nbr,dr_.nstatic) dr_.ghx zeros(endo_nbr,dr_.nfwr)];
B = dr_.ghu;
C = zeros(endo_nbr,(dr_.nstatic*(endo_nbr+1)));
for i=1:dr_.npred;
    C = [C dr_.ghxx(:,(i-1)*dr_.npred+1:i*dr_.npred)...
        zeros(endo_nbr,dr_.nstatic + dr_.nfwr)];
end
C = [C zeros(endo_nbr,dr_.nfwr*endo_nbr-dr_.nstatic)];
D = dr_.ghuu;
E = [zeros(endo_nbr,exo_nbr*dr_.nstatic) dr_.ghxu zeros(endo_nbr,exo_nbr*dr_.nfwr)];

%RATIONAL MODEL:
%=====
%Single technology shock in period 1 (NB productivity shock in 1st row,
% policy response/monetary shock in 2nd, wedge in 3rd):
shock = zeros(exo_nbr,t);

%SINGLE ONE SE PRODUCTIVITY SHOCK
%shock(2,1) = 0.0086;

%MULTIPLE RANDOM PRODUCTIVITY SHOCKS ~N(0,0.01)
shock(2,:) = reps;

```

```

%Calculate responses to one period technology shock in rational model
for i=2:t+1;
    X(:,i) = X(:,1)+A*(X(:,i-1)-X(:,1))+B*shock(:,i-1)+...
        0.5*C*kron(X(:,i-1)-X(:,1),X(:,i-1)-X(:,1))+...
        0.5*D*kron(shock(:,i-1),shock(:,i-1))+...
        E*kron(X(:,i-1)-X(:,1),shock(:,i-1));
end;

%BEHAVIOURAL MODEL:
%=====
chif = 0.5;
chic = 0.99/1.99;
%zetaq = 0.5;

%Single technology shock in period 1:
bshock = shock;
%causes repeated q shocks in future periods

%steady state values
BM1(:,1) = X(:,1); %variables from dynare model
ssq = X(strmatch('r_q', varlist, 'exact'),1);
ssrk = X(strmatch('r_rk', varlist, 'exact'),1);
BM2 = [ssq;ssq;ssq;ssq;ssq;ssq;0.5;0.5;ssrk;0;0;0;0;ssq;ssq];
%these are the behavioural variables:
behvars = {'bf_q';'b_q';'q(t-1)';'q(t-2)';'c_q';'cf_q';'f_q';'ff_q';...
    % 1 2 3 4 5 6 7 8
    'wc';'wf';'bf_rk';'zq';'eq';'zfq';'efq';'cf_q(t-1)';'ff_q(t-1)'};
% 9 10 11 12 13 14 15 16 17
% 1-8,14,15 in logs; 9,10 in levels; 11-13 log of the gross rate

for i=2:t+1
    %Policy response to previous period asset mispricing:
    if i>2;
    % bshock(4,i-1) = zetaq*(BM2(2,i-1)-ssq);
    % pol = zetaq*(q(t-1)-qss)
    bshock(4,i-1) = zetaq*(BM2(2,i-1)-...
        Y(strmatch('r_q', varlist, 'exact'),i-2));
    % pol = zetaq*(q(t-1)-qstar(t-1))
    % bshock(4,i-1) = zetaq*(BM2(2,i-1)-zetaq*(BM2(3,i-1)));
    % pol = zetaq*(q(t-1)-qstar(t-2))
    end;

    %first solve rationally to get qstar, rfrk, lambda and lambda(+1):
    BM1(:,i) = BM1(:,1)+A*(BM1(:,i-1)-BM1(:,1))+B*bshock(:,i-1)+...

```

```

0.5*C*kron(BM1(:,i-1)-BM1(:,1),BM1(:,i-1)-BM1(:,1))+...
0.5*D*kron(bshock(:,i-1),bshock(:,i-1))+...
E*kron(BM1(:,i-1)-BM1(:,1),bshock(:,i-1));

% bshock(4,i-1) = zetaq*(BM2(2,i-1)-BM1(strmatch('r_q', varlist, 'exact'),i));
% % pol = zetaq*(q(t-1)-qstar(t))
%
% %first solve rationally to get qstar, rfrk, lambda and lambda(+1):
BM1(:,i) = BM1(:,1)+A*(BM1(:,i-1)-BM1(:,1))+B*bshock(:,i-1)+...
0.5*C*kron(BM1(:,i-1)-BM1(:,1),BM1(:,i-1)-BM1(:,1))+...
0.5*D*kron(bshock(:,i-1),bshock(:,i-1))+...
E*kron(BM1(:,i-1)-BM1(:,1),bshock(:,i-1));

Y(:,i-1)=BM1(:,i);

%now solve behavioural model to get bfrk and wdg:
BM2(3,i) = BM2(2,i-1); % q(t-1)
BM2(4,i) = BM2(3,i-1); % q(t-2)
BM2(16,i) = BM2(6,i-1); % cf_q(t-1)
BM2(17,i) = BM2(8,i-1); % ff_q(t-1)
BM2(5,i) = log(exp(BM2(3,i))*(1+chic)+exp(BM2(5,i-1))*chic-2*chic*...
exp(BM2(4,i)));
% c_q = q(t-1)*(1+chic)+c_q(t-1)*chic-2*chic*q(t-2)
BM2(6,i) = log(exp(BM2(5,i))*(1+chic)-exp(BM2(5,i-1))*chic*(1+chic)+...
chic*(exp(BM2(6,i-1))-(1-chic)*exp(BM2(3,i))));
% cf_q = c_q*(1+chic)-c_q(t-1)*chic*(1+chic)+chic*(cf_q(t-1)-(1-chic)*q(t-1))
BM2(7,i) = log(exp(BM2(3,i))*(1-chif)+...
exp(BM1(strmatch('r_q', varlist, 'exact'),i))*chif);
% f_q = q(t-1)*(1-chif)+r_q*chif
BM2(8,i) = log(exp(BM2(7,i))*(1-chif)+...
exp(BM1(strmatch('rf_q', varlist, 'exact'),i))*chif);
% ff_q = f_q*(1-chif)+rf_q*chif
% BM2(9,i) = 0.5;
if i==2; BM2(9,i) = 0.5;
else if BM2(5,i-2)==BM2(7,i-2); BM2(9,i) = 0.5;
else profit(i) = (exp(BM2(2,i-1))-exp(BM2(2,i-2))*exp(BM1(strmatch('b_rb',...
varlist, 'exact'),i-2)));
c_profit(i) = profit(i)*sign(BM2(5,i-2)-BM2(4,i));
f_profit(i) = profit(i)*sign(BM2(7,i-2)-BM2(4,i));
BM2(9,i) = exp(3.75*c_profit(i))/(exp(3.75*c_profit(i))+exp(3.75*f_profit(i)))
end;
% [q(t-1)-q(t-2)*(1+rb(t-1))]
% clear profit c_profit f_profit;
end;
BM2(10,i) = 1-BM2(9,i);

```

```

    % wc/wf = wc(-1)/wf(-1)*exp(ff2(-2)-q(-1))^2/exp(cf2(-2)-q(-1))^2
    % wc+wf = 1
    BM2(1,i) = log(BM2(9,i)*exp(BM2(6,i))+BM2(10,i)*exp(BM2(8,i)));
    % bf_q = wc*cf_q+wf*ff_q;
    BM2(2,i) = log(BM2(9,i)*exp(BM2(5,i))+BM2(10,i)*exp(BM2(7,i)));
    % q = wc*c_q+wf*f_q;
    BM2(11,i) = exp(BM2(2,i)+BM1(strmatch('r_lambda', varlist, 'exact'),i)-...
        BM1(strmatch('rf_lambda', varlist, 'exact'),i))/beta-exp(BM2(1,i));
    % bf_rk = q*lambda/(beta*lambda(+1)) - bf_q
    BM2(12,i) = BM2(2,i)-BM1(strmatch('r_q', varlist, 'exact'),i);
    % zq = log(b_q/r_q)
    BM2(13,i) = BM2(12,i)-rhoq*BM2(12,i-1);
    % eq = zq - rhoq*zq(-1)
    BM2(14,i) = BM2(1,i)-BM1(strmatch('rf_q', varlist, 'exact'),i);
    % zfq = log(bf_q/rf_q);
    BM2(15,i) = BM2(14,i)-rhofq*BM2(14,i-1);
    % efq = zfq - rhofq*zfq(-1)

    %Add eq and efq 'shocks' into dynare model and resolve:
    bshock(3,i-1) = BM2(13,i);
    bshock(1,i-1) = BM2(15,i);
    BM1(:,i) = BM1(:,1)+A*(BM1(:,i-1)-BM1(:,1))+B*bshock(:,i-1)+...
        0.5*C*kron(BM1(:,i-1)-BM1(:,1),BM1(:,i-1)-BM1(:,1))+...
        0.5*D*kron(bshock(:,i-1),bshock(:,i-1))+...
        E*kron(BM1(:,i-1)-BM1(:,1),bshock(:,i-1));

    %Finally, put actual values into initial conditions for next period:
    actv = {'c';'i';'k';'l';'lambda';'mc';'mu';'pa';'pb';'pd';'pi';...
        'pstar';'q';'rb';'Rbi';'rk';'u';'v';'w';'wa';'wb';'wd';'wstar';'y'};
    for j=1:size(actv,1);
        BM1(strmatch(strcat('r_',actv(j,:)), varlist, 'exact'),i) = ...
            BM1(strmatch(strcat('b_',actv(j,:)), varlist, 'exact'),i);
    end;

end;

%Check consistency of bf_q, q, q(t-1), q(t-2), zq, zfq,
%across different matrices
temp = {'bf_q';'b_q';'zq';'zfq'};
for i=1:size(temp,1);
    if max(abs((BM1(strmatch(temp(i,:), varlist, 'exact'),:)-BM2(strmatch...
        (temp(i,:), behvars, 'exact'),:))))>.000001;
        warning('inconsistency across matrices');
        disp(temp(i));
    end;
end;

```

```

end;
if max(abs((BM1(strmatch('b_q', varlist, 'exact'),1:t)-BM2(strmatch...
('q(t-1)', behvars, 'exact'),2:t+1))))>.000001;
    warning('q(t-1) inconsistent across matrices');
end;
if max(abs((BM1(strmatch('b_q', varlist, 'exact'),1:t-1)-BM2(strmatch...
('q(t-2)', behvars, 'exact'),3:t+1))))>.000001;
    warning('q(t-2) inconsistent across matrices');
end;

%Compare utility
ratutility = 0.0;
for i=1:t-1;
    ratutility = ratutility+((beta^(i-1))*X(strmatch('r_v', varlist, 'exact'),i+1));
end;
ratutility = ratutility+((beta^(t-1))*X(strmatch('r_u', varlist, 'exact'),t+1));
%ratutility = (ratutility-X(strmatch('r_u', varlist, 'exact'),1))*100;
%disp(['utility in the rational model is equivalent to a '...
% num2str(ratutility) '% one-off change in consumption']);
behavutility = 0.0;
for i=1:t-1;
    behavutility = behavutility+((beta^(i-1))*BM1(strmatch('b_v', varlist, 'exact'),i+1));
end;
behavutility = behavutility+((beta^(t-1))*BM1(strmatch('b_u', varlist, 'exact'),t+1));
%behavutility = (behavutility-X(strmatch('r_u', varlist, 'exact'),1))*100;
%disp(['utility in the behavioural model is equivalent to a '...
% num2str(behavutility) '% one-off change in consumption']);
%
% %Print IRFs
% %Asset price
% figure ('papertype','A4','PaperUnits','centimeters','PaperPosition',...
% [0.5 0.5 20.5 27]);
% hold on;
% qplot = X(strmatch('r_q',varlist,'exact'),2:t+1)-X(strmatch('r_q',...
% varlist,'exact'),1);
% plot (qplot,'LineStyle','--','color','k');
% qplot = Y(strmatch('r_q',varlist,'exact'),1:t)-BM1(strmatch('r_q',...
% varlist,'exact'),1);
% plot (qplot,'color','g');
% qplot = BM2(strmatch('c_q',behvars,'exact'),2:t+1)-BM2(strmatch('c_q',...
% behvars,'exact'),1);
% plot (qplot,'color','b');
% qplot = BM2(strmatch('f_q',behvars,'exact'),2:t+1)-BM2(strmatch('f_q',...
% behvars,'exact'),1);
% plot (qplot,'color','c');

```

```

% qplot = BM1(strmatch('b_q',varlist,'exact'),2:t+1)-BM1(strmatch('b_q',...
%   varlist,'exact'),1);
% plot (qplot,'color','r');
% legend({'Rational Model','Rational expectation',...
%   "Chartist expectation","Fundamentalist expectation",...
%   'Actual'});
% title({'Asset Price'},'FontSize',10,'FontWeight','bold');
% hold off;
% print -dpdf IRF1.pdf;
% %Expectations of next period's asset price
% figure ('papertype','A4','PaperUnits','centimeters','PaperPosition',...
%   [0.5 0.5 20.5 27]);
% hold on;
% qplot = X(strmatch('rf_q',varlist,'exact'),2:t+1)-X(strmatch('rf_q',...
%   varlist,'exact'),1);
% plot (qplot,'LineStyle','--','color','k');
% qplot = BM1(strmatch('rf_q',varlist,'exact'),2:t+1)-BM1(strmatch('rf_q',...
%   varlist,'exact'),1);
% plot (qplot,'color','g');
% qplot = BM2(strmatch('cf_q',behvars,'exact'),2:t+1)-BM2(strmatch('cf_q',...
%   behvars,'exact'),1);
% plot (qplot,'color','b');
% qplot = BM2(strmatch('ff_q',behvars,'exact'),2:t+1)-BM2(strmatch('ff_q',...
%   behvars,'exact'),1);
% plot (qplot,'color','c');
% qplot = BM1(strmatch('bf_q',varlist,'exact'),2:t+1)-BM1(strmatch('bf_q',...
%   varlist,'exact'),1);
% plot (qplot,'color','r');
% legend({'Rational Model','Rational expectation',...
%   "Chartist expectation","Fundamentalist expectation",...
%   'Weighted average expectation'});
% title({'Expectations of the next period asset price'},'FontSize',10,...
%   'FontWeight','bold');
% hold off;
% print -dpdf IRF2.pdf;
% %Expectations of next period's return on capital
% figure ('papertype','A4','PaperUnits','centimeters','PaperPosition',...
%   [0.5 0.5 20.5 27]);
% hold on;
% qplot = X(strmatch('rf_rk',varlist,'exact'),2:t+1)-X(strmatch('rf_rk',...
%   varlist,'exact'),1);
% plot (qplot,'LineStyle','--','color','k');
% qplot = BM1(strmatch('rf_rk',varlist,'exact'),2:t+1)-BM1(strmatch('rf_rk',...
%   varlist,'exact'),1);
% plot (qplot,'color','g');

```

```

% qplot = BM1(strmatch('bf_rk',varlist,'exact'),2:t+1)-BM1(strmatch('bf_rk',...
%   varlist,'exact'),1);
% plot (qplot,'color','r');
% legend({'Rational Model','"Rational expectation"',...
%   '"Behavioural expectation"'});
% title({'Expectations of the next period return on capital'},'FontSize',...
%   10,'FontWeight','bold');
% hold off;
% print -dpdf IRF3.pdf;
% %Actual variables - rational model vs behavioural
% temp = {'c';'i';'k';'l';'y';'lambda';'mc';'mu';'pd';'pi';...
%   'pstar';'rb';'Rbi';'rk';'w';'wd';'wstar'};
% for i=1:2;
%   figure ('papertype','A4','PaperUnits','centimeters','PaperPosition',...
%     [0.5 0.5 20.5 27]);
%   for j=1:9;
%     subplot(3,3,j);
%     hold on;
%     if i*j==18; break; end;
%     title(strcat(temp(9*(i-1)+j,:),'_t'),'FontSize',10,...
%       'FontWeight','bold');
%     ratplot = X(strmatch(strcat('r_',temp(9*(i-1)+j,:)),varlist,...
%       'exact'),2:t+1)-X(strmatch(strcat('r_',temp(9*(i-1)+j,:)),...
%       varlist,'exact'),1);
%     plot (ratplot,'LineStyle','--','color','k');
%     behplot = BM1(strmatch(strcat('b_',temp(9*(i-1)+j,:)),varlist,...
%       'exact'),2:t+1)-BM1(strmatch(strcat('b_',temp(9*(i-1)+j,:)),...
%       varlist,'exact'),1);
%     plot (behplot,'color','r');
%     hold off;
%   end;
%   if i*j==18;
%     title('Weights (chartist dark)','FontSize',10,...
%       'FontWeight','bold');
%     plot (BM2(9,2:t+1),'LineStyle','--','color','b');
%     plot (BM2(10,2:t+1),'color','c');
%     hold off;
%   end;
%   eval(['print -dpdf IRF' int2str(i+3) '.pdf']);
% end;

```

B.6. Robustness Tests

Table B.1. Utility cost of various Taylor rules in the model with no switching between forecasting rules

Difference compared to benchmark model: $v = 0$		
ζ_Q	Utility cost versus rational model (%age permanent consumption)	Performance relative to $\zeta_Q = 0$ (%age bias corrected)
0.00	-0.184 %	0.0 %
0.05	-0.197 %	-6.6 %
-0.05	-0.173 %	6.2 %
0.10	-0.211 %	-14.6 %
-0.10	-0.163 %	11.8 %
0.50	-0.412 %	-123.2 %
-0.50	-0.111 %	39.9 %
1.00	-1.642 %	-790.5 %
-1.00	-0.101 %	45.5 %

Table B.2. Utility cost of various Taylor rules in the model with rapid switching between forecasting rules

Difference compared to benchmark model: $v = 100$		
ζ_Q	Utility cost versus rational model (%age permanent consumption)	Performance relative to $\zeta_Q = 0$ (%age bias corrected)
0.00	-0.233 %	0.0 %
0.05	-0.186 %	20.2 %
-0.05	-0.293 %	-25.6 %
0.10	-0.165 %	29.5 %
-0.10	-0.358 %	-53.4 %
0.50	-3.815 %	-1534.4 %
-0.50	-0.595 %	-154.9 %
1.00	NA %	NA %
-1.00	-0.895 %	-283.6 %

Table B.3. Utility cost of various Taylor rules in the model with a 1 percent uncertainty bound around the fundamentalist forecast

Difference compared to benchmark model: $C = 0.01$		
ζ_Q	Utility cost versus rational model (%age permanent consumption)	Performance relative to $\zeta_Q = 0$ (%age bias corrected)
0.00	-0.163 %	0.0 %
0.05	-0.176 %	-7.9 %
-0.05	-0.153 %	6.4 %
0.10	-0.190 %	-16.3 %
-0.10	-0.143 %	12.5 %
0.50	-0.373 %	-129.1 %
-0.50	-0.097 %	40.2 %
1.00	NA %	NA %
-1.00	-0.114 %	30.3 %

Table B.4. Utility cost of various Taylor rules in the model with a 10 percent uncertainty bound around the fundamentalist forecast

Difference compared to benchmark model: $C = 0.1$		
ζ_Q	Utility cost versus rational model (%age permanent consumption)	Performance relative to $\zeta_Q = 0$ (%age bias corrected)
0.00	-0.130 %	0.0 %
0.05	-0.190 %	-46.0 %
-0.05	-0.090 %	30.5 %
0.10	-0.285 %	-119.2 %
-0.10	-0.063 %	51.5 %
0.50	NA %	NA %
-0.50	-0.004 %	97.0 %
1.00	NA %	NA %
-1.00	-0.059 %	54.4 %