

**Essays in Financial Economics: Option Pricing,
Behavioural Finance, Stochastic Terms in
Modelling, and Relationships between Sectors
within Stock Markets**

Konstantina Mari

PhD

University of York

Economics and Related Studies

March 2017

Abstract

Finance Theory is built mainly based on the assumption that investors are risk neutral. We run two finance experiments to test American option pricing theory. Dynamic programming and optimal stopping theory are used for the construction of the models. Simulation studies are conducted for the experiments. In the first experiment, we test a disinvestment case of real options theory in discrete time. Our results show that risk aversion explains better the decisions of the participants than risk neutrality as few of the participants appeared to be risk neutral. Furthermore, risk aversion explains better the behaviour of the subjects than myopic behaviour. In the second experiment, we examine the timing of the exercising of an American call option contract in continuous time. We estimate the risk aversion parameters of the subjects from the main experiment, and we elicit their risk aversion parameters by running another small experiment with allocation questions. Based on these parameters we find the estimated and the elicited risk averse optimal trigger respectively. From our analysis we show that the estimated risk averse optimal trigger explains better the actual stopping decisions of the subjects, while the risk neutral optimal trigger has the next highest explanatory power and the elicited risk optimal trigger the lowest. The third project tests the stochastic assumptions underlying an analysis by examining the statistical properties of the estimated parameters and comparing them to the actual values. This study shows that the stochastic specification underlying any analysis matters for the interpretation of its results. In the last project, we check the interdependency relationships among the five major market sectors of Greece, Italy and Portugal. By using dependency tests, such as Johansen cointegration and Granger causality tests, we find that the Greek sectors provide some diversification benefits for sector-level investments, while the opposite is true for the Italian sectors. Only in the case of Portugal do the results suggest non-existence of interdependency relationships among the sectors before the financial crisis and their existence after it. Moreover, by using the variance decomposition and the time-varying volatility methodologies we conclude that for all the three countries the majority of the sectors are exogenous and their volatility is highly increased after the financial crisis, particularly in the case of the Financials sector.

To my family

“Είμαστε αυτό που πράττουμε επανειλημμένα. Έτσι, η τελειότητα δεν είναι πράξη αλλά συνήθεια.”

Αριστοτέλης

“We are what we repeatedly do. Excellence then is not an act but a habit. “

Aristotle

Contents

Abstract	3
Table of Contents	7
List of Figures	9
List of Tables	10
Acknowledgements	12
Declaration	14
1. Introduction	15
2. Do People Disinvest Optimally?	19
2.1 Introduction	19
2.2 Theory	21
2.3 Experimental implementation	25
2.4 Analysis and results	27
2.5 Conclusions	30
Appendix A	41
Appendix B	49
3. Optimal Timing of the Exercising of a Financial Option Contract	58
3.1 Introduction	58
3.2 Theory	60
3.2.1 The significance of geometric Brownian motion in asset pricing	61
3.2.2 Optimal timing of the exercising of a financial option contract	65
3.3 Experimental design and implementation	72
3.3.1 Experimental design	72
3.3.2 Experimental implementation	77
3.4 Analysis and results	82
3.4.1 Relation of actual subjects' trigger with risk-neutral and risk-averse optimal trigger	84
3.4.2 Relation of actual subjects' trigger with risk aversion and other variables	88
3.4.3 Relation of actual subjects' trigger with risk-neutral and risk-averse optimal trigger including the data for the rounds in which the subject had not exercised by the time the option expired.....	93
3.5 Conclusions	97
Appendix A	101
Appendix B	103
Appendix C	114
4. Does the Stochastic Specification Matter?	118

4.1 Introduction	118
4.2 The allocation method for eliciting risk-aversion	119
4.3 The optimal allocation	121
4.4 Assumed stochastic specifications	122
4.4.1 Additive normal (an)	122
4.4.2 Beta (<i>b</i>).....	123
4.4.3 Random preferences normal (<i>rpn</i>)	123
4.4.4 Additive logistic (al).....	124
4.4.5 Random preferences lognormal (<i>rpl</i>)	124
4.5 The simulation and estimation program	124
4.6 Results	126
4.7 Conclusions	129
5. Stock Market Sectors' Dependencies Pre and Post the Financial Crisis of 2007: the cases of Greece, Italy and Portugal	146
5.1 Introduction	146
5.2 Data	149
5.3 Methodology.....	150
5.3.1 Dickey-Fuller GLS unit root test	152
5.3.2 The Johansen cointegration test.....	154
5.3.3 Granger causality test	156
5.3.4 Generalised impulse responses and variance decomposition	158
5.3.5 Dynamic conditional correlation (DCC)	160
5.3.6 Time-varying volatility	162
5.4 Empirical results	163
5.4.1 Greece.....	163
5.4.2 Italy	172
5.4.3 Portugal.....	181
5.5 Summary and conclusions.....	190
Appendix A	193
Appendix B	194
Appendix C	195
Appendix D	195
Appendix E.....	196
Appendix F.....	197
Appendix G	202
Appendix H.....	207
References.....	212

List of Figures

Chapter 2	
1. An Example of a Binomial Tree	38
2. Subject 66	39
3. Histogram of Maximum Percentage Consistent Over All 74 Subjects	40
<i>Additional Figures in the Instructions</i>	
Chapter 3	
1. Graph related to the Derivation of the Option Value formula	70
2. Theory Graph of Trigger and Exercise Price formulas	72
3. A Screen Shot from the First Part Software	79
4. A Screen Shot from the Second Part Software	81
5. Scatter of Actual Trigger against Risk Neutral Optimal Trigger	86
6. Scatter of Actual Trigger against Elicited Risk Averse Optimal Trigger	87
7. Scatter of Actual Trigger against Estimated Risk Averse Optimal Trigger	87
8. Scatter of Actual Trigger against Risk Neutral Optimal Trigger.....	95
9. Scatter of Actual Trigger against Elicited Risk Averse Optimal Trigger.....	95
10. Scatter of Actual Trigger against Estimated Risk Averse Optimal Trigger.....	96
<i>Additional Figures in the Instructions</i>	
Chapter 4	
1. Beta distributions of x for $x^*=0.05$ (green on the left), 0.25 red (in the middle) and 0.5 blue (on the right); $s=50$	132
2. Normal (blue dashed) and logistic (red solid) distributions	132
3. Normal (blue dashed) and lognormal (red solid) distributions	132
4. Simulated allocations based on preference parameters 1 and problem number 29	133
Chapter 5	
<i>Figures in the Appendices:</i>	
App. F.1 Impulse responses – Greece	197
App. F.3 Time-varying volatilities and correlations – Greece	201
App. G.1 Impulse responses – Italy	202
App. G.3 Time-varying volatilities and correlations – Italy	206
App. H.1 Impulse responses – Portugal	207
App. H.3 Time-varying volatilities and correlations – Portugal	211

List of Tables

Chapter 2

1. The Problem Set	31
2. The 'Winning' hypotheses subject by subject	32
3. A Classification of the Winners	35
4. Tests of Significance	36

Chapter 3

1. Part 1 – the Problem Set	75
2. Practice and Payment Rounds	78
3. Regression of the actual trigger on the risk-neutral optimal trigger	84
4. Regression of the actual trigger on the risk-neutral optimal trigger (without intercept)	85
5. Regression of the actual trigger on the elicited risk-neutral optimal trigger	85
6. Regression of the actual trigger on the estimated risk-averse optimal trigger ...	86
7. Regression of the actual trigger on the estimated risk-averse optimal trigger and gender	88
8. Regression of the actual trigger on the estimated risk-averse optimal trigger and all the final statistically significant demographic variables	90
9. Regression of the actual trigger on the elicited risk-averse optimal trigger and all the statistically significant demographic variables	92
10. Truncated regression of the actual trigger on the risk-neutral optimal trigger...	93
11. Truncated regression of the actual trigger on the elicited risk-averse optimal trigger.....	94
12. Truncated regression of the actual trigger on the estimated risk-averse optimal trigger.....	94
T1: Problems Set for the Allocation Decision Problem Part	100

Chapter 4

1. Maximised Log-Likelihoods	134
2. r estimates	137
3. Standard deviation of r estimates	140
4. s estimates	143

Chapter 5

1. Descriptive statistics of returns in different periods – Greece	165
2. Dickey – Fuller GLS unit root test – Greece	166
3. Multivariate Johansen cointegration test – Greece	167
4. Granger causality test results – Greece	168
5. Descriptive statistics of returns in different periods – Italy	173
6. Dickey – Fuller GLS unit root test – Italy	175
7. Multivariate Johansen cointegration test – Italy	176

8. Granger causality test results – Italy	178
9. Descriptive statistics of returns in different periods – Portugal	182
10. Dickey – Fuller GLS unit root test – Portugal	185
11. Multivariate Johansen cointegration test – Portugal	186
12. Granger causality test results – Portugal	187
<i>Additional Tables in the Appendices:</i>	
App. F.2 Variance decomposition – Greece.....	199
App. G.2 Variance decomposition – Italy.....	204
App. H.2 Variance decomposition – Portugal.....	209

Acknowledgements

This thesis is result of hard work, stressful days and nights, long time continuous dedication only on research and continuous willingness to exceed the “known” and learn more and more. I would like to thank a lot my supervisor, Professor John Hey, for all the intellectual discussions that we had and for all his very helpful advices on how this thesis to look as it is today. John, I will keep in my mind all the significant moments during our meetings, your way of thinking in research, and the exchange of ideas that we had during all these years. I would also like to thank Dr Vanessa Smith for being supportive and a good advisor in my last chapter, a new area in my PhD thesis that I was glad to work on. Finally, I also thank the rest of my Thesis Advisory Panel members, Professor Peter Simmons and Mr. John Bone, for all their positive view on my work and all their comments.

Having completed this thesis, I would like to thank all the Institutions and Foundations which supported me financially during my PhD studies. Specifically, I would like to thank the Foundation for Education and European Culture for the scholarship they offered me for conducting my postgraduate studies, the RIS (Research and Impact Support, University of York) for their contribution in my research expenses, and the Department of Economics and Related Studies for the Teaching scholarships that I received during the previous years and the financial grants that it provided me for the presentation of my work at International Conferences. Furthermore, I would also like to thank the Royal Economic Society for providing the funds to attend an Easter school.

During these long days and nights (up to morning hours sometimes) in the PhD area, I had the chance to have the most wonderful breaks; some very interesting and intellectual discussions with my colleagues. In addition, funny moments did not miss too. I would like to thank all these people for all the nice academic discussions that we had and the cheerful

times that we passed in restaurants and pubs. Apart from my friends in the PhD area, I want also to thank my friends back in Greece who were always supportive and close to me, my dreams and my worries.

Thankful for the knowledge I obtained during the previous years, I can continue happy now my journey in the academic world as an Associate Lecturer at Queen Mary University of London. This thesis is dedicated to my family who was supporting me from the beginning up to the end of this journey. Without your continuous love, support, help, encouragement and dedication nothing would have been the same. Thank you for everything you gave me and for always being there for me.

Declaration

I, Konstantina Mari, declare that this thesis titled, “Essays in Financial Economics: Option Pricing, Behavioural Finance, Stochastic Terms in Modelling, and Relationships between Sectors within Stock Markets” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made reference to that.
- This work has not previously been presented for an award at this, or any other, University.

Chapter 1

Introduction

Finance theory is built based on the assumption that investors are rational in a very strong sense, and, often, that they are risk-neutral. However, empirical and experimental studies have proved that economic behaviour is much more complex. Even if all the investors try to maximise their utility, some of them deviate from the optimal choice due a number of reasons, such as irrationality, inefficiency or liquidity constraints. Moreover, some of these studies have proved that preference-based explanations play a key role in financial markets and explain better the behaviour of people than risk neutrality. More specifically, Levy et al (2000) argue that one of the explanations that dominates in markets is risk aversion. Therefore, it is rational to check if this theory can explain better than risk neutrality the actual behaviour of people in the context of option pricing theory. As option contracts are some of the most used derivative assets in financial markets, it is crucial to check the reliability of option pricing theory in real world. In the second chapter and the third chapter of the thesis, we test experimentally the American option pricing theory and more specifically the optimal exercising time of this type of contract. We also test risk aversion as a reason that causes people to deviate from the optimal decision.

The disinvestment decision is of crucial importance in many contexts: if funds are tied up for too long in a poorly-performing project, then opportunities for re-investment may be missed. Optimal disinvestment theory is a component of real options theory; in some senses it is simply the converse of optimal investment theory, but is relatively ignored by

experimentalists. Two papers that do consider it conclude that decision-makers stay in projects longer than that prescribed by the optimal behaviour of a risk-neutral agent. This departure is 'explained' through risk-aversion, but without a formal hypothesis under test. In the second chapter, we report on a similar experiment that we run, but we explain the behaviour of the subjects with the estimation of the level of risk-aversion. We also explore an alternative hypothesis to explain behaviour – that subjects are myopic. Our results show that few subjects appear to be risk-neutral, many seem to be risk-averse but few are myopic. In the third chapter, we test one fundamental financial theory in continuous time, that of the optimal time of exercising an American call option contract. We notice that few experimental papers test continuous time theories and that there is only one other paper that tests experimentally this theory but that paper focuses on the learning process of the subjects. Its authors suspect that risk aversion is the reason that the subjects deviate from the risk neutral optimal trigger but they do not test it formally. Through a two-stage experiment that we run, we check if risk neutrality or risk aversion explains better the behaviour of the subjects. More specifically, we test if the estimated risk averse optimal trigger, which is extracted by the first stage (main experiment), or the elicited risk averse optimal trigger, which is elicited by some allocation questions (second stage of the experiment), explains better than the risk neutral optimal trigger the actual decision of the subjects. In our analysis, we check also the explanatory power of each of the two risk averse optimal triggers in conjunction with some demographic variables on the actual behaviour of the subjects. Our conclusions are that the estimated risk averse optimal trigger does better than the elicited risk averse optimal trigger generally, while the risk neutral optimal trigger explains the actual choices of the subjects better than the elicited risk averse one but worse than the estimated risk averse optimal trigger.

The fourth chapter would appear to take us off at a tangent, and in some senses this is true, but it highlights the importance and significance of the stochastic assumptions underlying any empirical analysis. The context is not either of the contexts of the chapters 2 and 3, but the message that emerges from chapter 4 is an important one: that the stochastic specification underlying any statistical analysis matters for the interpretation of its results. Usually the stochastic assumptions in any data analysis are implicit, rather than explicitly, stated. For example, in chapter 2 it was implicitly assumed that the observations were independent observations of a binary random variable, while in chapter 3 it was assumed that the divergence between the actual triggers and the theoretically optimal ones could be described by a zero-mean homoscedastic normally distributed random variable (this is the usual assumption lurking behind any regression analysis). Chapter 4 considers a different context: the statistical properties of stated allocations (in an allocation problem) and their relation to the optimal allocation. It is closest to the analysis contained in chapter 3. Our results show that whether the stochastic specification matters depend on what one is interested in. Chapter 4 suggests that before doing any statistical analysis one should carry out extensive simulations. This is an important message - though perhaps not one that we ourselves have listened to in the other chapters of the thesis. Still one cannot do everything in a single thesis!

The fifth chapter investigates the interdependency links among the stock market sectors of Greece, Italy and Portugal before and after the global financial crisis of 2008. The Johansen cointegration and the Granger causality tests show that Greek market sectors provide good diversification benefits as the interdependency relationships are very limited due to the lack of cointegration and strong causality relationships among them in both pre-crisis and post-crisis periods. On the contrary, the same tests indicate that the diversification benefits from the Italian sectors are limited when it comes to sector-level investments because of the strong appearance of long-run and short-run relationships among the sectors for both the

pre-crisis and post-crisis periods. Portuguese sectors are not cointegrated before the financial crisis with just very few causality relationships, but they are found to be cointegrated with more causality relationships after it. The variance decomposition results indicate that most of the sectors are exogenous both before and after the financial crisis with the Financials sector being the most endogenous out of all across the three countries. Finally, through our investigation it is found that the time-varying volatilities of the sectors are shown to have substantially increased after the financial crisis.

Chapter 2

Do People Disinvest Optimally?¹

2.1 Introduction

This chapter reports on an experimental study of the disinvestment problem. In this problem, the decision-maker (DM) holds an asset which yields stochastic cash flows until its disposal. There is a deadline for the disposal of the asset. The decision problem consists of deciding when to dispose of the asset; the optimal disposal point is dependent on both the time period and the cash flow at that time².

Clearly this problem is a special case of the class of dynamic decision problems, and, more particularly, of the class of real options problems. It is in some ways the converse of the optimal investment problem, but we single out its investigation because of its importance in many fields of economics, affecting the performance of many firms and individuals. Moreover, although theorists may regard it as the converse of an investment problem, it remains to be seen whether actual decision-makers regard it as such.

Our inspiration for this study are the papers by Sandri *et al* (2010) and by Musshoff *et al* (2013), who experimentally explore a disinvestment problem. Their main experimental finding is implicit in the title of their articles, and is that many subjects hold on to the asset for longer than that prescribed by the theory appropriate for a risk-neutral DM. They

¹ This chapter is based on join work with John Hey.

² And perhaps other things – depending upon the objective function – as we shall show later.

conjecture from this that risk-aversion may have a role to play. The way that they model this is not to assume that the objective function of the DM is a concave function of the payoff(s), but by using “risk-adjusted discount rates”. Their analysis concludes with their Proposition P3: “The larger an individual’s risk aversion, the earlier the disinvestment occurs” in the first paper and the similar Hypothesis H4: “Risk-averse farmers disinvest earlier” in the second. These propositions seem to go against the conclusions of the theoretical paper by Henderson and Hobson (2013), in which they report that risk-aversion *may* delay the disinvestment decision: it depends upon the objective function.

Sandri *et al* and Musshoff *et al* elicited risk-attitudes independently of the disinvestment problem, using a Holt-Laury (2002) price list, and used these elicited risk-attitudes to explain observed behaviour. In contrast, in this paper, we fit to the data two models of risk-averse DMs, and find the risk-attitude and the model which best explains the behaviour of our subjects. Thus, our risk indices are not elicited independently of the disinvestment problem, but estimated from the disinvestment behaviour. The reason for this is that it would not be clear what one could infer from the results if the independent elicitation differed from that implied by behaviour – other than that the elicitation method influenced the elicitation result. Furthermore, it is proved³ that different elicitation methods of the risk aversion parameter give different results. Therefore, this study shows that the level of risk aversion of a person is different from decision task to decision task. Hence, it seems that the best way to find the risk aversion coefficient as an explanatory variable of a person’s decisions in a financial risky problem is based on the person’s decisions in this specific decision task.

We also explore an alternative hypothesis in an attempt to explain behaviour. This hypothesis embodies the idea that DMs are unable to solve the backward induction solution

³ For example, in the paper of Loomes and Pogrebna (2014).

to the dynamic problem, and instead use a shorter horizon, which they roll forward period-by-period as the problem unfolds. We call this the Rolling Strategy; we give details later.

The chapter is structured as follows. In the next section we describe the decision problem and find its optimal solution for two different objective functions (both embodying risk-aversion); we also find its solution for a DM who follows the Rolling Strategy. In section 3 we describe the experimental implementation, and in section 4 we present our findings. Section 5 concludes.

2.2 Theory

We start from the set-up of Sandri *et al.* We operate in a discrete world. The DM owns some asset which must be disposed of by some final period, denoted by T . Until the asset is disposed of, the DM receives cash flows every period. These cash flows follow a binomial random walk: if we denote the cash flow in period t by x_t , then the cash flow in period $t+1$ is either x_t+h or x_t-h with respective probabilities p and $1-p$. The parameters h and p are constant. The theory usually embodies a discount rate applied to future incomes, but, because of the impossibility of having real discounting in an experiment lasting around one hour, we introduce interest on the disposal value of the asset, L , from the time when disposal occurs until T . So if the asset is disposed of in period t the value of the disposal is Lr^{T-t} where r is the rate of return (one plus the rate of interest). Interest is not received on the cash flows. After disposal takes place no further cash flows are received. The DM receives the value of the asset plus interest and the sum of the cash flows experienced until disposal. We have the usual trade-off: the later disposal takes place the more cash flows are received but the lower the disposal value of the asset.

The solution is found in the normal way, using backward induction. We start with a risk-neutral DM and later generalise to a risk-averse DM. Let us denote by $V_t(x_t)$ the expected payoff to the DM as viewed from period t when the cash flow in that period is x_t . In the final period T the DM must dispose of the asset if he or she has not done so earlier. It follows that

$$V_T(x_T) = L + x_T \quad (1)$$

Let us proceed to the general backward induction on V . In any period t , the DM disposes of the asset if the payoff from so doing exceed the expected payoff from continuing to hold the asset. So we have

$$V_t(x_t) = \max[x_t + Lr^{T-t}, pV_{t+1}(x_t+h) + (1-p)V_{t+1}(x_t-h)] \quad (2)$$

the first term in this expression being the payoff from disposing of it now and the second term the expected payoff from continuing to hold the asset. The decision in t is implicit in this expression: if the first term is the maximum it is best to dispose of it now; if the second term is the maximum it is best to continue holding it. Notice that previous cash flows do not enter into this equation as they cancel out on both sides. This provides the optimal strategy for a risk-neutral DM.

Now we turn to a non-risk-neutral DM, that is, one with a non-linear function $u(\cdot)$ over payoffs. In many dynamic decision problems, with the monetary payoff in period t denoted by y_t , the objective function is normally assumed to be the maximisation of the expectation of the expression

$$u(y_1) + u(y_2) + \dots + u(y_T) \quad (3)$$

Taking this to be the objective function, which we are going to call Objective Function 1 (OF1), equations (1) and (2) above become

$$V_T(x_T) = u(L + x_T) \quad (1')$$

and

$$V_t(x_t) = \max[u(x_t + Lr^{T-t}), pV_{t+1}(x_t+h)+(1-p)V_{t+1}(x_t-h)] \quad (2')$$

As before the sum of the utilities of previous cash flows cancel out from both sides of (2').

This provides the optimal strategy for a DM with OF1, which is given by equation (3). Obviously the solution depends upon the form of the utility function $u(\cdot)$. In what follows we consider both the Constant Relative Risk Aversion (CRRA) form and the Constant Absolute Risk Aversion (CARA) form.

However, in the context of an experiment OF1 may seem a bit odd. The subject walks out of the laboratory with some money; presumably therefore it is the expected value of the utility of that money which concerns him or her. If that is so, then the objective is the maximisation of the expectation of the expression

$$u(y_1 + y_2 + \dots + y_T) \quad (4)$$

This we call Objective Function 2 (OF2) – the maximisation of the expected value of the utility of the sum of the payoffs – in contrast to OF1 – the maximisation of the expected value of the sum of the utilities of the payoffs.

It may appear that OF2 is more reasonable from a behavioural point of view – though it is only by looking at behaviour can we be sure; in the analysis section we shall see which fits behaviour best. For the moment we note a complication with computing the optimal strategy. With OF2 past cash flows no longer cancel out of the two sides of equation (2'). This means that the decision in any period with any given cash flow in that period depends *also* on the accumulated cash flows to that point. Equation (2') is no longer valid. We present the solution in Appendix A.

Finally, we describe in more detail the Rolling Strategy. This embodies the notion that the DM is myopic and looks ahead every period only S periods; the DM has a short horizon that is rolled forward every period. If $S=T-1$ then the DM behaves as above, but if $S < T-1$ then until

period $T-S+1$ the DM acts as if he or she thinks that he or she has to dispose of the asset in period $t+S$. So if $S=3$ and $T=6$, then in period 1 the DM acts as if he or she thinks that he or she has to dispose of the asset in period 4, in period 2 the DM acts as if he or she thinks that he or she has to dispose of the asset in period 5, while in periods 3, 4, 5 and 6 the DM correctly acts as if he or she thinks that he or she has to dispose of the asset in period 6. Obviously it is not optimal, but it is not clear in general how much the DM loses by using it.

We assume that a DM who uses the Rolling Strategy is risk-neutral⁴. The implied decision rules can be immediately found from those above, though some new notation is required. Let us denote by $D(t, x_t, T)$ the optimal decision (either 1 for continue or 0 for disposal) of a risk-neutral DM (implied by equation (2)) in time period t with cash flow x_t when the true horizon is T , and denote by $d(t, x_t, S, T)$ the decisions of someone with a rolling horizon of S in the same position but with a true horizon of T .

We have that:

-if $t \geq T-S$ then $d(t, x_t, S, T) = D(t, x_t, T)$ because the true horizon is within the correct horizon

-if $t < T-S$ then $d(t, x_t, S, T) = D(t, x_t, t+S)$ because the DM is optimising under the (wrong) assumption that he/she *has* to liquidate in period $t+S$.

Full details are given in Appendix A.

⁴ Alternatively we could have assumed risk-aversion, so that all the models considered are nested within this model.

2.3 Experimental implementation

The experiment was carried out in the EXEC laboratory. Subjects started by reading the written Instructions⁵; they could ask questions for clarification. The experiment was computerised, with the code written in Visual Studio. The core of the experiment was a binomial tree – a screen shot is in Figure 1. Subjects were presented with a series of 15 different problems (chosen after extensive simulation by us⁶). All problems had a final disposal period of $T=15$. At the start of each problem they were told: the initial cash flow, x_1 ; the disposal value, L ; the rate of interest on the disposal value, r ; the jump size, h ; and the probability of an upward jump, p . They then had to decide each period (until they disposed of the asset) whether to dispose of it or not. If they chose not to dispose of it, Nature moved randomly (according to the specified probability) and the appropriate part of the tree eliminated. Figure 1 shows the situation in a problem with initial cash flow 20, liquidation value 280, interest rate 5%, jump size 5 and probability of jumping up 0.9, where the DM has decided not to dispose of the asset in the first period, Nature has moved Up and it is now the time for the DM to decide what to do in the second period. In order to encourage the subjects to think about the problem the ‘Continue’ and ‘Stop’ buttons did not appear until 20 seconds had elapsed, but they were restricted to a total time of 40 seconds in any one period⁷. The ‘Confirm’ button did not appear until they had clicked on either ‘Continue’ or ‘Stop’.

In the tree, each vertically-aligned pair of small boxes represents a node that the DM might reach (depending on their decisions and Nature’s moves) – with the top number indicating the cash flow at that node and the bottom number indicating the probability of getting to

⁵ The Instructions are available in Appendix B.

⁶ The problems were chosen in such a way that different degrees of risk aversion, or different rolling horizons, would imply different optimal strategies – so that we could distinguish between subjects. See the [simulation code](#).

⁷ If they had not taken a decision by the time that these 40 seconds had elapsed, the software assumed that the subject wished *not* to dispose of the asset.

that node (if the DM had not disposed of the asset earlier). So, in the example in Figure 1, where the DM is in period 2, the possible nodes reachable in period 3 imply cash flows of 30, if Nature moves Up, or 20, if Nature moves Down (with respective probabilities 0.9 and 0.1). At the left-hand side of the screen (not shown in Figure 1) was a summary of the Instructions, and, as will be seen from the figure, information is provided about the move that Nature has just made, the cash flow of that period, and the disposal value of the asset if disposed of in that period (the disposal value plus interest to the end of the problem). Written Instructions were also provided, with a number of examples. Subjects read the Instructions before the experiment started. Any questions were answered by the experimenters.

Subjects were paid the sum of the show up fee (£2.50) and their total payoff in one of the 15 problems chosen randomly, using the exchange rate of 100 tokens equalling £1. There was a total of 74 subjects; the average payment was £10.67.

We note that the 15 problems were chosen after extensive simulations using Matlab following a small pilot study. A key requirement for the problem set was that we could infer from the data the level of risk-aversion of the subjects, and/or the length of their rolling horizon (if they were following the rolling strategy). This meant that we needed problems where different levels of risk-aversion, or different lengths of the rolling horizon, implied different decisions. We should briefly explore what this means and how we took it into account.

The decisions that we observe are binary decisions: either Stop (coded 0) or Continue (coded 1). So a strategy is defined by a set of 0's or 1's at each of the nodes in the tree. Let us confine the discussion to OF1⁸. The nodes in a tree of length T consist of 1 in the first period, 2 in the second period, ..., t in the t 'th period, ..., up to $T-1$ in the $(T-1)$ 'th period: a total of $(T-1)T/2$

⁸ With OF2 we need to define nodes not just by the period and the vertical positioning in the tree at that point but also by the accumulated cash flow at that point.

nodes⁹. At each of these nodes a strategy implies a decision D which is either 0 or 1. If we add together these 0's and 1's we get a total (call it N) which indicates the number of continuation nodes in the tree. We want different values of N for different levels of risk aversion (or of the rolling horizon). So we looked for problems where the variance of N over a reasonable set of risk aversion levels (or over different horizons) was as large as possible. We also wanted problems where a risk-neutral person behaving optimally would expect to earn around £10 (the conversion rate between experimental tokens and money was 1 token = 1p). To decide on our problem set we carried out pre-experimental simulations. The problem set is given in Table 1.

An alternative experimental task, possibly for future work, would have been to ask subjects to specify a *strategy* – a statement of what they would do at each possible decision node that they might reach. Given the number of such decision nodes this would have been time-consuming and complicated. We could simplify the task of specifying a strategy by asking them to impose a *threshold* in each time period – above which they would liquidate – but this would suggest to them the nature of the optimal strategy, and we would be pushing them towards the 'correct' solution.

2.4 Analysis and results

The data that we have on each subject and in each problem are the decisions of the subject at each node that the subject reached – for each problem, the data are a string of 1's followed by a '0'. Because of the nature of the data, we proceed as follows. We would like to know, for each of OF1 and OF2, for each utility function (CRRA and CARA), which level of risk-aversion, and for the Rolling strategy which horizon, best explains the behaviour of each

⁹ We ignore the final period as the compulsory decision is to dispose of the asset.

subject, and how much of that behaviour it explains. Suppose a subject takes a total of N decisions throughout the course of the experiment (this number varies from subject to subject), then, for each Objective Function, for each utility function (CRRA and CARA), and for each level of risk-aversion, and for the Rolling strategy for each horizon, we can calculate how many of the decisions of the subject are consistent with that specification. We can then calculate the percentage of the decisions taken by the subject that are consistent with that specification. Call this pc . It depends in general on the specification. This variable definitely is not smooth; if we graph it against, for example, the risk aversion, or against the rolling horizon, it is a step function. Finding its maximum (which would give us the best-fitting risk-aversion or the best-fitting horizon) using Maximum Likelihood techniques would not work due to the fact that it is not smoothly concave. But we need to find where it reaches its maximum and find its value at the maximum. We show this graphically; Figure 2 illustrates – this is for subject 66. On the vertical axis is the variable pc . On the horizontal axis is the rolling horizon S for the rolling strategy. The horizontal axis also indicates the level of risk-aversion, with risk-aversion *decreasing* from left to right, going from very risk-averse at the left to risk-neutrality at the right when $x=14$. There are five ‘curves’ in the picture, with – indicating OF1 CRRA and so on. It will be seen that all the five curves finish (at $x=14$) at the same vertical point. This is because when the horizon is the correct horizon (14) and the DM is risk-neutral, all five strategies lead to the same decisions. So for this subject, assuming he or she is risk-neutral, or working with the correct horizon, explains just 70% of the subject’s decisions. It will be noticed that all the curves are indeed step-functions: the horizontal line shows that for OF2 CARA changing the level of risk-aversion has no effect on decisions; this is an implication of the CRRA utility function in the context of OF2. The greatest percentage correct is with OF1 CRRA – with 91% correct; the second highest is with OF1 CARA – with 79% correct. We might ask whether 91% is significantly larger than 79%. This, of course, depends upon the number of decisions, which, for this subject, was 43. Carrying out the

standard test as to whether one proportion is greater than another shows that 91% is significantly greater than 70% at 10% but not at 5%. However the hypothesis that the subject is risk-averse explains significantly more than the hypothesis that the subject is risk-neutral. The CRRA utility function is $u(x)=x^r$ and the best-fitting value of r for this subject is between¹⁰ 0.70 and 0.73 – a moderately risk-averse person.

Table 2 lists the best-fitting specifications and the best-fitting risk-aversion index or best-fitting rolling horizon subject by subject. For some subjects there are clear unique winners – as Table 3 shows. If Risk-Neutrality is the best, then, since all the other specifications have risk-neutrality nested within them, we do not list them for the other specifications. So a subject whose behaviour is listed in the category OF1 CRRA is *strictly* risk-averse, and so on. We note that OF2 CRRA only appears the best for the risk-neutral subjects; in fact, as a glance at Figure 2 will show, the optimal decisions for someone with an OF2 CRRA preference functional are *not* dependent on the level of risk-aversion¹¹, and are therefore the same as for a risk-neutral subject.

It is clear from Table 3 that the Rolling Strategy does not do well – only coming joint winner with OF1 CRRA for two subjects. The best-performing specification is OF1 CRRA coming first on its own for 38 subjects, and joint first with OF1 CARA 10 times.

It is of interest to see whether the winning specification is *significantly* better than the others, and, in particular, significantly better than risk-neutrality. Table 4 gives the details of standard t-tests of the difference between two proportions. It will be seen that with the exception of the 5 risk-neutral subjects, for 43 out of our 74 subjects the winning

¹⁰ We only get a range estimate because the pc function is horizontal at its peak – again a consequence of the the data that we have.

¹¹ This is a consequence of the CRRA utility function.

specification fits significantly better than risk-neutrality at 1%, for 6 subjects at 5%, and for 1 subject at 10%.

2.5 Conclusions

It is clear from these results that Sandri *et al* and Musshoff *et al* were right – risk-aversion is needed to rationalise the behaviour of the subjects; our Rolling Strategy does not do very well. It is also clear from our results that risk-aversion varies considerably across the subjects. Taking into account the risk-aversion significantly improves the results – as Table 4 shows. Looking at Figure 3 it seems that we can explain much of the behaviour with one or other of our specifications for many of our subjects.

However, Figure 3 also shows that for some subjects we can only rationalise a rather small percentage of their decisions (as low as 62% for one subject). This suggests that there may be some other decision rule that these subjects were following. One possibility is that subjects thought that Nature may have had a memory (though this was not true). In this case, earlier sequences of moves by Nature may have affected their future decisions. Another heuristic could be that subjects try to avoid any possible future regret which can occur if they continue following the moves of Nature or they decide based on the regret that they have already experienced by their past decisions and the released moves by Nature. Also, the gambler's fallacy is a cognitive bias which could possibly explain the subjects' behaviour as people tend to believe in this kind of problems that if an outcome (random event) happens many times in a sequence during a period, it will happen less frequently in the future, and vice versa. Investigating these cases may be of interest for future work. After all, backward induction is a complicated and computationally complex procedure. It would not be surprising if subjects developed simple heuristics for tackling the problem.

Tables

Table 1: The Problem Set

Number	p	h	L	r	x_1
1	0.6	1	40	1.25	10
2	0.1	1	75	1.2	20
3	0.3	1	40	1.25	20
4	0.8	1	150	1.15	15
5	0.5	1	75	1.2	20
6	0.2	2	80	1.2	20
7	0.3	2	135	1.15	35
8	0.9	5	280	1.05	20
9	0.7	1	80	1.2	5
10	0.4	2	140	1.15	30
11	0.8	1	45	1.25	0
12	0.1	2	40	1.25	25
13	0.6	1	70	1.2	20
14	0.3	1	45	1.25	10
15	0.4	5	260	1.1	5

p : probability of jumping Up

h : jump size

L : disposal value

r : rate of return on disposal

x_1 : initial cash flow

Table 2: The ‘Winning’ hypotheses subject by subject (Note: ranges are given when the best-fitting value is not unique)

SUBJE CT	OF1 CRRA	OF1 CARA	OF2 CRRA	OF2 CARA	ROLLING	RN
1	0.11-0.41	0.024-0.0445				
2	0.83-1	0-0.001	0.11-1	0-0.018	7-14	Y
3	0.11-0.5	0.0095-0.0445				
4	0.66 , 0.7					
5				0.0185		
6	0.68-0.69					
7	0.58-0.59					
8	0.11-0.5	0.03-0.0445				
9	0.58					
10				0.029		
11	0.11-0.5 , 0.54-0.58	0.0065 , 0.0075-0.0445				
12	0.68-0.69					
13	0.11-0.5	0.007-0.008 , 0.031-0.0445				
14	0.59					
15				0.0185		
16				0.031 , 0.0355		
17	0.7-0.72					
18	0.58-0.59 , 0.62-0.63 , 0.69					
19	0.67-0.68					
20				0.0185		
21	0.81-0.82				5	
22	0.67					
23				0.0355 , 0.0365		
24				0.029		
25	0.67					
26	0.7					
27	0.67 , 0.7-0.71					

28	0.58-.059					
29	0.58					
30		0.005				
31	0.67-0.69					
32	0.67 , 0.7					
33	0.71-0.74					
34	0.6-0.63					
35	0.75 , 0.79-0.82 , 0.84-1	0-0.002	0.11-1	0-0.018	7-14	Y
36	0.81-0.82					
37	0.59					
38	0.66					
39	0.6-0.63					
40	0.11-0.46	0.011-0.0445				
41				0.0395		
42				0.039		
43	0.59					
44				0.0285-0.029 , 0.0355 , 0.0365 , 0.0385		
45	0.83-1	0-0.001	0.11-1	0-0.018	7-14	Y
46	0.73-0.75					
47				0.019		
48	0.73-0.75					
49	0.11-0.5 , 0.54-0.59	0.0075-0.0445				
50	0.66				2 , 3	
51	0.77-0.8	0.0015-0.002				
52	0.11-0.5 , 0.54-0.55 , 0.58-0.59	0.006-0.0065 , 0.0075-0.008 , 0.031-0.0445				
53	0.11-0.5	0.005 , 0.007-0.0445				
54		0.013-0.0305				
55				0.019, 0.0205		
56	0.6-0.63 , 0.67					
57	0.59 , 0.67					

58	0.67					
59				0.019		
60	0.67-0.69					
61	0.58-0.59 , 0.62-0.63 , 0.65 , 0.67-0.69 , 0.71	0.004		0.0445		
62	0.67 , 0.7-0.71					
63	0.58-0.59 , 0.62-0.63 , 0.68					
64				0.019		
65				0.0185		
66	0.7-0.73					
67				0.0185		
68	0.83-1	0-0.001	0.11-1	0-0.018	7-14	Y
69	0.83-1	0-0.001	0.11-1	0-0.018	7-14	Y
70	0.11-0.5	0.0085-0.0445				
71	0.59					
72				0.0365 , 0.0385		
73	0.68 , 0.7-0.72 , 0.76					
74	0.6-0.63					

Table 3: A Classification of the Winners

specification(s)	<i>n</i>
RN	5
OF1 CRRA	38
OF1 CARA	2
OF2 CARA	16
OF1 CRRA <i>and</i> OF1 CARA	10
OF1 CRRA <i>and</i> Rolling	2
OF1 CRRA, OF1 CARA <i>and</i> OF2 CARA	1
Total	74

n = the number of subjects for which this/these are the best

Table 4: Tests of Significance

subject	Number decisions	p	Risk-Neutral?	t-stat 1 st vs 2 nd	t-stat 1 st vs 3 rd	t-stat 1 st vs 4 th	t-stat 1 st vs RN
1	129	88	N	0,00	3,04***	6,16***	11,42***
2	29	100	Y	-	-	-	-
3	68	78	N	0,00	1,19	3,73***	4,51***
4	44	77	N	0,95	0,95	1,14	1,14
5	21	90	N	0,40	0,40	0,40	0,40
6	50	84	N	0,76	0,76	2,86***	2,86***
7	73	82	N	0,16	2,20**	3,05***	5,43***
8	77	83	N	0,00	3,76***	5,14***	7,46***
9	52	79	N	0,25	0,48	0,94	2,50***
10	65	80	N	0,42	0,42	2,26**	4,44***
11	65	80	N	0,00	0,28	2,49***	3,80***
12	50	90	N	1,40*	1,64**	2,90***	3,65***
13	44	68	N	0,00	0,69	1,90**	2,73***
14	51	78	N	0,36	1,24	1,45*	2,26**
15	22	73	N	0,36	0,36	0,36	0,36
16	51	75	N	0,23	0,45	1,92**	4,06***
17	31	77	N	1,04	1,36*	2,06**	2,36***
18	61	80	N	0,14	1,39*	1,97**	3,89***
19	88	90	N	0,82	2,17**	2,76***	7,19***
20	37	81	N	0,32	0,32	0,32	0,32
21	29	72	N	0,00	0,25	0,25	0,49
22	56	82	N	0,40	0,40	0,66	3,38***
23	43	77	N	0,32	0,74	1,03	1,03
24	35	69	N	0,53	0,53	0,79	1,70**
25	53	83	N	1,13	2,01**	2,22**	3,12***
26	39	79	N	0,21	0,52	0,52	0,52
27	52	85	N	1,04	1,04	1,94**	2,85***
28	83	89	N	0,20	3,80***	3,92***	6,48***
29	50	74	N	0,23	1,49*	3,24***	4,40***
30	68	82	N	0,15	0,44	0,44	7,00***
31	45	73	N	0,21	0,21	0,62	0,92
32	51	82	N	0,50	0,86	2,15**	2,35***
33	50	86	N	0,55	1,04	2,93***	2,93***
34	62	81	N	0,28	0,55	2,01**	3,94***
35	38	76	Y	0,00	0,00	0,00	0,00
36	48	90	N	0,31	1,49*	2,14**	3,39***
37	65	82	N	0,29	1,10	1,10	4,38***
38	50	78	N	0,47	0,47	1,54*	2,14**
39	58	83	N	0,55	1,18	2,85***	3,56***
40	98	86	N	0,00	2,41***	2,99***	7,94***
41	75	84	N	0,16	0,48	2,79***	5,66***

42	49	76	N	0,56	0,56	1,19	2,57***
43	74	84	N	0,32	1,49*	2,77***	5,63***
44	38	68	N	0,46	0,46	1,60*	1,85**
45	35	83	Y	0,00	0,00	0,00	0,00
46	24	67	N	0,64	0,64	0,64	0,64
47	25	72	N	0,31	0,31	0,31	0,61
48	26	65	N	0,52	0,81	1,09	1,09
49	83	86	N	0,00	1,64**	3,27***	6,49***
50	31	71	N	0,00	0,26	1,54*	1,84**
51	39	85	N	0,00	0,69	0,90	1,20
52	60	80	N	0,00	1,50*	2,92***	3,65***
53	71	82	N	0,00	1,02	2,30**	5,68***
54	146	95	N	0,37	4,43***	5,46***	12,51***
55	22	68	N	0,28	0,28	0,28	0,28
56	55	76	N	0,36	1,26	1,26	4,94***
57	54	78	N	0,49	1,28*	1,71**	2,63***
58	58	81	N	0,40	1,72**	3,00***	4,52***
59	39	79	N	0,52	0,52	0,52	0,52
60	46	78	N	0,66	0,66	2,61***	2,61***
61	54	80	N	0,00	0,00	0,74	2,87***
62	45	82	N	0,47	0,70	0,70	1,02
63	46	72	N	0,21	1,50*	1,79**	1,79**
64	25	68	N	0,59	0,59	0,59	0,59
65	32	69	N	0,26	0,26	0,26	0,26
66	43	91	N	1.56 *	2.07 **	2.46 ***	2.46 ***
67	24	79	N	0,33	0,33	0,33	0,33
68	35	83	Y	0,00	0,00	0,00	0,00
69	18	78	Y	0,00	0,00	0,00	0,00
70	84	87	N	0,00	2,68***	5,75***	6,68***
71	59	81	N	0,40	1,51*	1,85**	2,71***
72	45	78	N	0,97	0,97	1,17	1,46*
73	29	62	N	0,23	0,23	1,07	1,07
74	50	76	N	0,46	1,31*	1,51*	4,22***

p : the percentage of decisions consistent with the best-fitting specification.

First entry is the test statistic, the asterisks indicate significance:

* at 10% (1,28), ** at 5% (1,64), *** at 1% (2,32)

Figures

Figure 1: An Example of a Binomial Tree

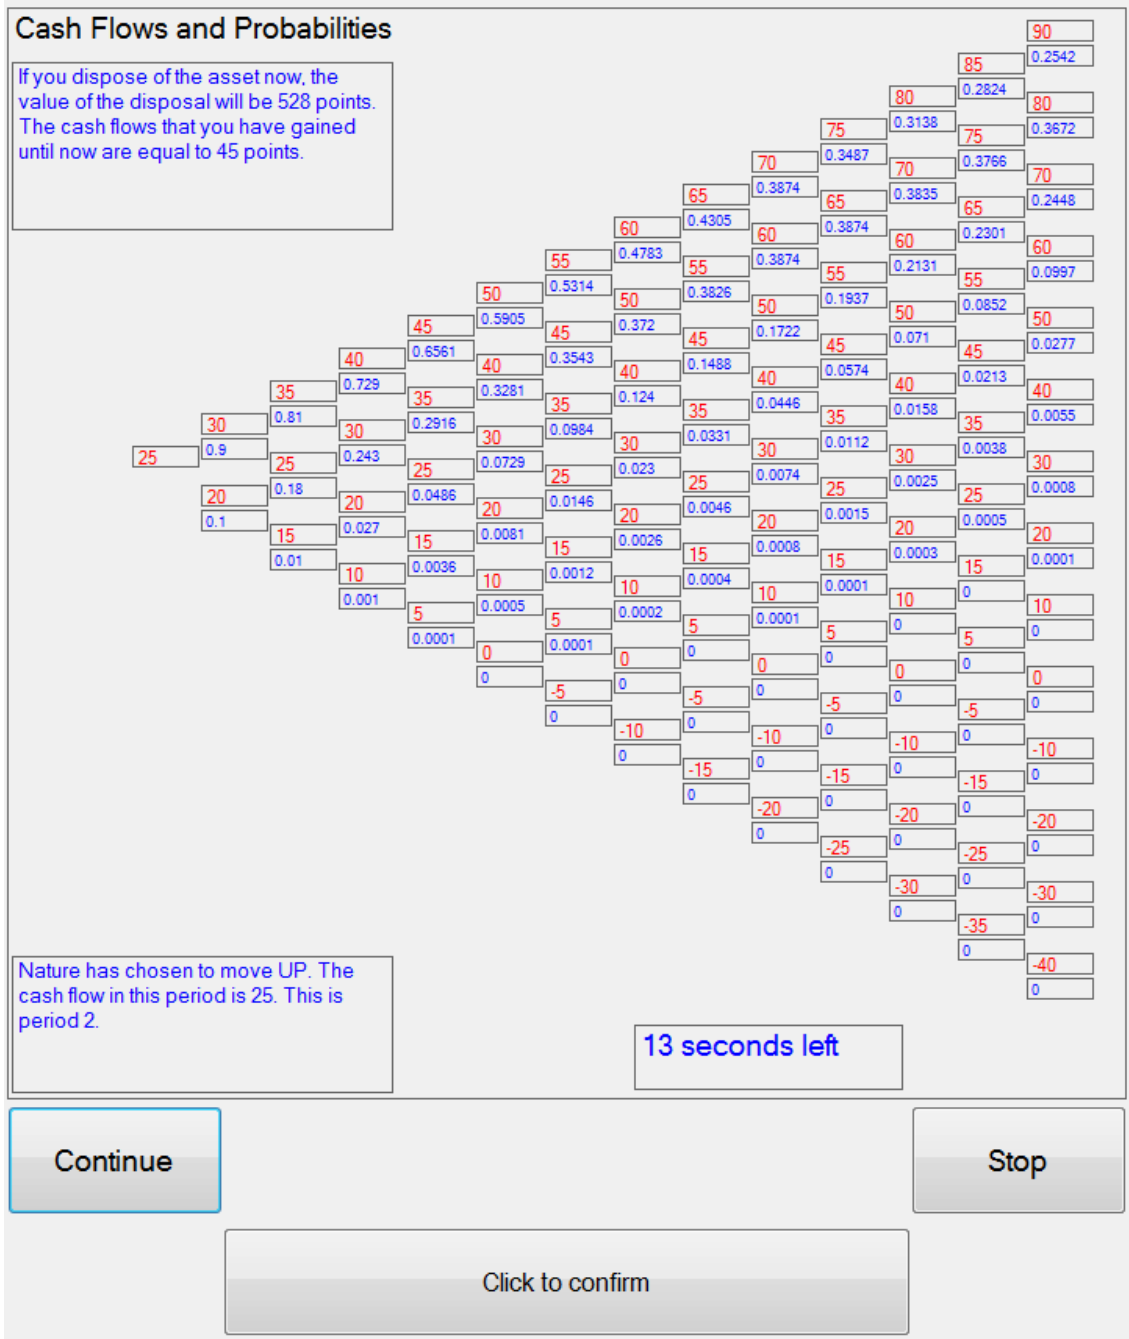




Figure 2: Subject 66



OF1 CRRA —
OF1 CARA --
OF2 CRRA x
OF2 CARA .
RH —

At some points the 'curves' overlap.

Figure 3: Histogram of Maximum Percentage Consistent Over All 74 Subjects



Appendix A

Theory and Programming

The DM owns an asset with liquidation value L which earns interest at the rate of return rr until the forced liquidation date in period T . We start in period 1 with a cash flow x_1 . Thereafter the cash flow follows a binomial random walk: given x in period t the cash flow in period $t+1$ is either $x+h$ (with probability p) or $x-h$ (with probability $1-p$).

We start with **Objective Function 1**: the maximisation of the expectation of the sum of the utilities of the payoffs (cash flow plus liquidation when it occurs).

There are various *nodes* that the DM may reach. In period 1 there is just 1; in period 2 there are 2;...; in period t , there are t of them;..., in period T there are T of them. The total number of such nodes is $1+2+\dots+T = T(T+1)/2$. We could refer to these nodes with a *pair* of numbers (t,j) where j goes from 1 to t in period t . A better way is to define k nodes which go sequentially from 1 to $T(T+1)/2$. In period 1, k is just 1; in period 2, k is 2 and 3; in period t , k goes from $(t-1)t/2+1$ to $t(t+1)/2$; in period T , k goes from $(T-1)T/2+1$ to $T(T+1)/2$. We call the total number of k nodes *totk*. This is equal to $T(T+1)/2$.

At each k node there is an associated cash flow. Letting x_k denote the cash flow at the node k . From the binomial process we have the following Matlab code:

```
for t=2:1:T % going through the periods
  for j=1:1:t % for each period going through the j nodes
    k=(t-1)*t/2+j; % calculating the corresponding k node
    x(k)=x(1)+(t-2*j+1)*h; % calculating the cash flow at that k node
  end
end
```

Now let us find the solution for Objective Function 1, where the objective is the maximisation of the expected value of the sum of the utilities. We use the following notation. d_k is the optimal decision at node k . EV_k is the expected value of the objective function *as viewed from*

node k . At this node the previous elements $u(x_1) + u(x_2) + \dots$ are given and known and therefore do not enter the objective function.

Denote by lqv_k the liquidation value of liquidating at that node, and by ctv_k the continuation (expected) value at that node.

From node k the DM either moves Up or Down. We need to know to which k nodes these moves take us. From the tree (see Figure 1) it can be seen that if at node k in period t moving Up takes the DM to node $k+t$ in period $t+1$, while moving Down takes the DM to node $k+t+1$ in period $t+1$.

We work backwards starting in period T . Here there are no decisions to take and we have simply for k between $(T-1)T/2+1$ and $T(T+1)/2$ that:

$$(1) \quad EV_k = u(x_{k+L}) \quad \text{This is for the period } T \text{ nodes.}$$

We work backwards now. Here we take the general case of $k \leq (T-1)T/2$ (that is in periods 1 through $T-1$). We first write the solution in equations and then transfer it into Matlab code. The backward induction starts in period $T-1$ and then works backwards to period 1. In period t (note that in period t the index k takes values from $(t-1)t/2+1$ to $t(t+1)/2$ inclusive) the relevant equations are:

$$(2) \quad ctv_k = u(x_k) + pEV_{k+t} + (1-p)EV_{k+t+1}$$

$$(3) \quad lqv_k = u(x_k + Lrr^{T-t})$$

$$(4) \quad d_k = 1 \text{ if } ctv_k \geq lqv_k; 0 \text{ otherwise (We are assuming a DM who is indifferent continues).}$$

$$(5) \quad EV_k = \max[ctv_k, lqv_k]$$

The Matlab code follows (note we use here a generic utility function; in the code we distinguish between CRRA and CARA).

```

for t=T-1:-1:1 % for each period working back
for k=1+t*(t-1)/2:1:t*(t+1)/2 % for the k nodes in that period
ctv(k)=u(x(k))+p*EV(k+t)+(1-p)*EV(k+t+1); % continuation value
lqv(k)=u(L*(rr^(T-t))+x(k)); % liquidation value
if lqv(k)<=ctv(k) % if continuing is better
EV(k)=ctv(k); % the continuation value is EV
d(k)=1; % decision is to continue
end
if lqv(k)>ctv(k) % if liquidating is better
EV(k)=lqv(k); % the liquidation value is EV
d(k)=0; % the decision is to liquidate
end
end
end
end

```

Now let us turn to **Objective Function 2**, where the objective function is the maximisation of the expected utility of the sum of payoffs. This means that the optimal decision at any k -node depends not only on that node but also the accumulated cash flows at that node. Note crucially that knowing one is at a particular k -node is not sufficient to know the accumulated cash flow at that node; this latter depends upon the *route* by which the DM has reached that node. For example consider $k=5$ in $t=3$. This node could have been reached by going Up from period 1 to 2 and then Down from period 2 to 3; or it could have been reached by going Down from period 1 to 2 and then Up from period 2 to 3. In the former case the accumulated cash flow would be $x_1 + (x_1+h) + x_1 = 3x_1+h$; in the latter case the accumulated cash flow would be $x_1 + (x_1-h) + x_1 = 3x_1-h$. In order to deal with this, we need to introduce what we call l -nodes, indicating not only which k -node the DM is at, but also the accumulated cash flow he or she has. We should note that two different l -nodes do not necessarily have different accumulated cash flows.

How many l -nodes are there? It can be seen from Figure 1 that in period t there are a total of 2^{t-1} l -nodes, half of them reached by going Up from the 2^{t-2} l -nodes in period $t-1$ and half of them reached by going Down from the 2^{t-2} l -nodes in period $t-1$. Thus, the total number of l -nodes in a tree of length T is $1 + 2 + 2^2 + 2^3 + 2^4 + \dots + 2^{T-1} = 2^T - 1$. We need to calculate the

optimal decisions at all of these with the exception of the 2^{T-1} nodes in period T where the only decision is to stop. Hence we have (where the subscript now is the l -node):

$$(6) \quad d_l = 0 \text{ if } 2^{T-2} + 1 \leq l \leq 2^{T-1}.$$

Also we have (in the final period if it is reached):

$$(7) \quad EV_l = u(X_l + L) \text{ if } 2^{T-2} + 1 \leq l \leq 2^{T-1} \text{ where } X_l \text{ denotes the accumulated cash flow at node } l, \text{ and } EV_l \text{ denotes the value of the objective function at node } l.$$

Now the optimisation procedure is straightforward. We already have the (default) decisions in the final period and the Expected Value of the objective function at each of the final period l -nodes. Hence we can write (recall that the vector $upl(l)$ tells us to which l -node a movement Up from node l takes the DM, and $dnl(l)$ tells us to which l -node a movement Down takes the DM from node l):

For all the other l -nodes in periods $t < T$ we have:

$$(8) \quad ctv_l = pEV_{upl(l)} + (1-p)EV_{dnl(l)}$$

$$(9) \quad lqv_l = u(X_l + Lrr^{T-t})$$

$$(10) \quad d_l = 1 \text{ if } ctv_l \geq lqv_l; 0 \text{ otherwise (We are assuming a DM who is indifferent continues).}$$

$$(11) \quad EV_l = \max[ctv_l, lqv_l]$$

Note that in these expressions the value of t is that corresponding to the period in which that l -node is in.

Let us number the l nodes so that we have the following.

<i>t</i>	<i>k</i>	<i>e(k) the number of <i>l</i> nodes in the <i>k</i> node</i>	<i>l</i>	up / node	in <i>k</i> node	down /- node	in <i>k</i> node
1	1	1	1	2	2	3	3
2	2	1	2	4	4	6	4
	3	1	3	5	5	7	6
3	4	1	4	8	7	11	8
	5	2	5	9	8	13	9
			6	10	8	14	9
6	1	7	7	12	9	15	10
4	7	1	8	16	11	20	12
	8	3	9	17	12	24	13
			10	18	12	25	13
			11	19	12	26	13
	9	3	12	21	13	28	14
			13	22	13	29	14
14			23	13	30	14	
10	1	15	27	14	31	15	
5	11	1	16	32	16	37	17
	12	4	17	33	17	44	18
			18	34	17	45	18
			19	35	17	46	18
			20	36	17	47	18
	13	6	21	38	18	52	19
			22	39	18	53	19
			23	40	18	54	19
			24	41	18	55	19
			25	42	18	56	19
			26	43	18	57	19
	14	4	27	48	19	59	20
			28	49	19	60	20
			29	50	19	61	20
			30	51	19	62	20
15	1	31	58	20	63	21	

Hence the implied tree and the *j*, *k* and *l* nodes are as follows.

1 (1)	2 (2 to 3)	3 (4 to 7)	4 (8 to 15)	5 (16 to 31)	6 (32 to 63)
					1; 16; 32
				1; 11; 16	
			1; 7; 8		2; 17; 33, 34, 35, 36, 37
		1; 4; 4		2; 12; 17, 18, 19, 20	
	1; 2; 2		2; 8; 9, 10, 11		3; 18; 38, 39, 40, 41, 42, 43, 44, 45, 46, 47
1; 1; 1		2; 5; 5, 6		3; 13; 21, 22, 23, 24, 25, 26	
	2; 3; 3		3; 9; 12, 13, 14		4; 19; 48, 49, 50, 51, 52, 53, 54, 55, 56, 57
		3; 6; 7		4; 14; 27, 28, 29, 30	
			4; 10; 15		5; 20; 58, 59, 60, 61, 62
				5; 15; 31	
					6; 21; 63

The numbers in the top row are the *t*-values.

At each node, the first number is what we call the j -value; the second number is the k -node; and all the other numbers are the l -nodes. Of these other numbers, the ones in normal font are the l -nodes reached by coming DOWN from the previous period, and those in *italics* those reached by coming UP.

Thus, if we number the l nodes this way, it is nice and simple: l goes up to $2l$ and goes down to $2l+1$, for all l from 1 to $2^{T-2}-1$ (up to the penultimate period)

Moreover the accumulated cash flow at node l is the cash flow in the associated k node plus the accumulated cash flow in the k node from where it came.

$$(12) \quad X(l) = X(l/2) + x(k) \text{ if } l \text{ is even}$$

$$(13) \quad X(l) = X((l-1)/2) + x(k) \text{ if } l \text{ is odd}$$

where k is the k node in which l is.

We should do this for all l from 2 to 2^{T-1} .

Now how to find the k node in which a particular l value is. The following Matlab code appears to work. Note that there are $(t-1)!/[(j-1)!(t-j)!]$ l nodes in node (t,j) . This expression is calculated using `nchoosek` in Matlab.

```
l=0;
k=0;
clk(1)=1;
for t=2:1:T
    for j=1:1:t
        k=k+1;
        n1=nchoosek(t-1,j-1);
        for ll=1:1:n1
            l=l+1;
            clk(l)=k;
        end
    end
end
```

We also need to know (see above for the liquidation values) the t node corresponding to a particular k node. Here is the Matlab code the vector `ckt`:

```

% now we need to find the t value corresponding to any k value
k=0;
for t=1:1:T
    for j=1:1:t
        k=k+1;
        ckt(k)=t;
    end
end

% this does the important stuff
for l=2^(T-1):1:2^T-1; % these are the period T nodes
    if rt==1
        EV(l)=crra(L+X(l),r);
    end
    if rt==2
        EV(l)=cara(L+X(l),r);
    end
end

% this is the important recursion for the other periods going backwards
% this is for a generic utility function
for t=T-1:-1:1
    for l=2^(t-1):1:2^t-1
        kk=clk(l);
        tt=ckt(kk);
        ctv(l)=p*EV(2*l)+(1-p)*EV(2*l+1);
        lqv(l)=u(L*(rr^(T-tt))+X(l));
        if lqv(l)<=ctv(l)
            EV(l)=ctv(l);
            d(l)=1;
        end
        if lqv(l)>ctv(l)
            EV(l)=lqv(l);
            d(l)=0;
        end
    end
end
end

```

Finally let us show the decisions of a DM with a **rolling strategy**. Here we assume risk-neutrality.

We need to start with the fully-optimal strategy – backwardly inducting from the end. Let us denote the *Expected Value* to the decision-maker of fully optimising if he or she is at node k by $EV_{T,k}$ (the first argument indicating the horizon used by the decision-maker and the second the node). Let us denote by $D_{T,k}$ the optimal decision, taking the value 1 for continuing and the value 0 for liquidating. We work backwards. We are now making the notation consistent with the Matlab code.

In T we have (ignoring the accumulated cash flows which are given and the decision-maker will get anyhow):

$$(14) \quad D_{T,k} = 0 \quad \text{for } k \text{ from } 1+(T-1)T/2 \text{ to } T(T+1)/2$$

$$(15) \quad EV_{T,k} = x_k + L \quad \text{for } k \text{ from } 1+(T-1)T/2 \text{ to } T(T+1)/2$$

Now we work *backwards*, from $t=T-1$ to $t=1$, using the following recursion. Note that if the DM is at node k in period t , then going up arrives at node $k+t$ in period $t+1$, and going down arrives at node $k+t+1$ in period $t+1$.

$$(16) \quad D_{T,k} = 0 \text{ if } x_k + Lr^{(T-t)} > pEV_{T,k+t} + (1-p)EV_{T,k+t+1}$$

$$= 1 \text{ if } x_k + Lr^{(T-t)} \leq pEV_{T,k+t} + (1-p)EV_{T,k+t+1}$$

$$(17) \quad EV_{T,k} = \max[x_k + Lr^{(T-t)}, x_{tj} + pEV_{T,k+t} + (1-p)EV_{T,k+t+1}]$$

So we have the optimal decision at each cash flow node.

Now let us consider someone who has a rolling strategy with an horizon of S periods – so in period t works *as if* he or she *has* to liquidate in period $t+S$ or in period T whichever is the sooner (the true liquidation date is T). Let us use $d_{S,T,k}$ to denote the decision of such a decision-maker at node k , the first argument indicating the rolling horizon, the second the true horizon and the third the node.

Be careful about the notation: $D_{T,k}$ denotes the *optimal* decision at node k for an optimising decision who has to liquidate in period T . In contrast $d_{S,T,k}$ denotes the decision at node k of a DM with a rolling horizon of S periods ahead in a problem where he/she actually has to liquidate in period T but wrongly working on the presumption that they have to liquidate S periods ahead.

It follows that we have the following results:

$$(18) \quad \text{If } t \geq T-S \text{ then } d_{S,T,k} = D_{T,k} \text{ because the true horizon is within the correct horizon.}$$

$$(19) \quad \text{If } t < T-S \text{ then } d_{S,T,k} = D_{t+S,k} \text{ because the DM is optimising under the (wrong) assumption that he/she } \textit{has} \text{ to liquidate in period } t+S.$$

Appendix B

Experimental Instructions

The Instructions of the Experiment which were given to the subjects follow.



Instructions

Preamble

Welcome to this experiment. Thank you for coming. You are going to participate in an experiment. Please read carefully the instructions, they are to help you to understand what you will be asked to do. You are going to earn money for your participation in the experiment and you will be paid immediately after the completion of the experiment.

The Experiment

You will be presented with a sequence of 15 different problems. In each problem you start owning some asset which you have to dispose of during or at the end of the problem. It has a value. When you dispose of the asset, this value will earn interest until the end of the problem, and this will constitute part of your payment for that problem. In each problem there is a sequence of time periods and you can dispose of the asset in any of these. In each

of the time periods, until you dispose of the asset you will also earn a cash flow (which could be positive or negative) and which will be added to (or subtracted from if it is negative) your payment from the disposal. The cash flow follows a random path, determined by Nature, jumping either up or down by a fixed specified amount each period with specified and fixed probabilities. The disposal decision can be taken in any one of the periods of the problem, though you will have to dispose of it in the final period if you have not disposed of it before then. More specifically, your task is to decide in each time period in each of the problems whether you want to continue holding the asset to the next period of the problem, or whether you want to stop and dispose of the asset in that period. You will be given the following information in each problem: the value of the asset on disposal, the initial cash flow, the size of the jump in the cash flow, the probability that the cash flow jumps up (and the residual probability that the cash flows move down), the number of periods in the problem at the end of which you have to dispose of it, and the rate of interest on the disposal value.

The Interface of the Experiment

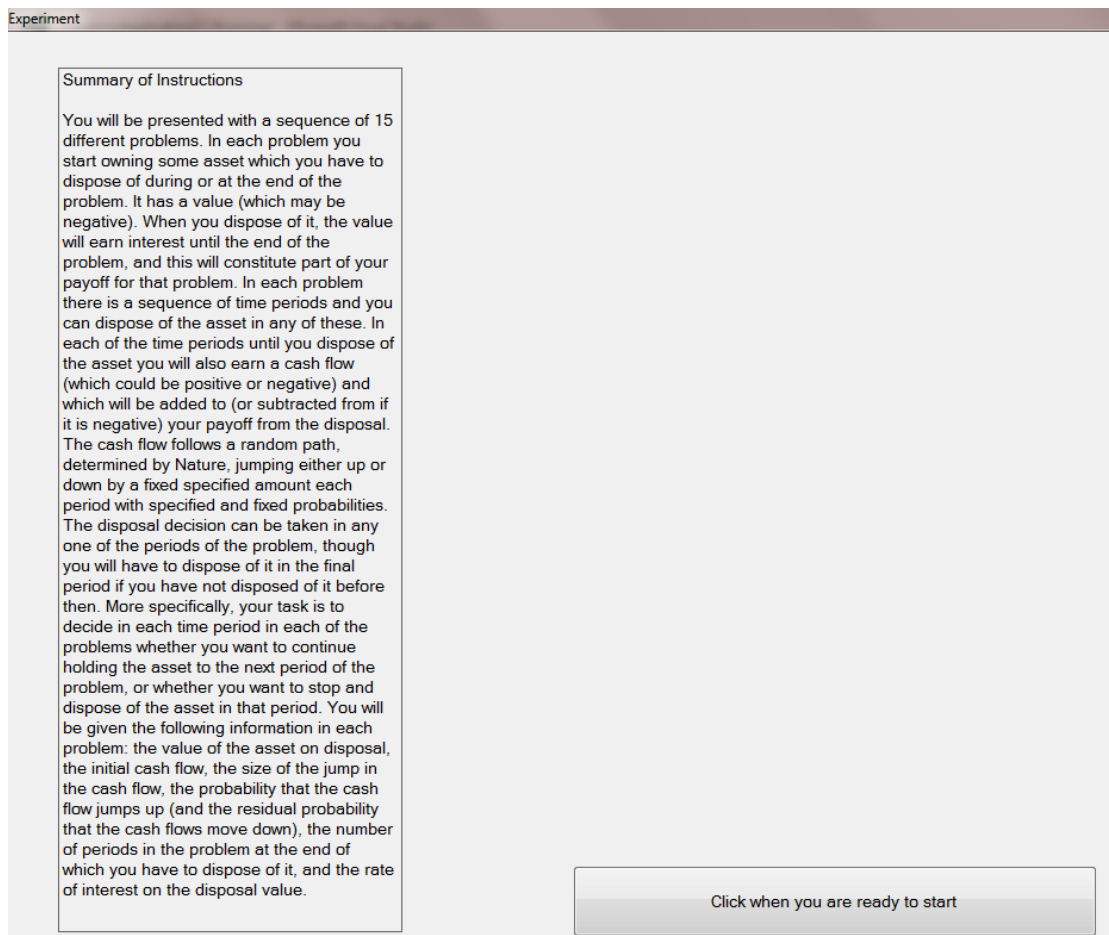


Figure 1

At the beginning of the experiment you will see the Introduction to these Instructions. After reading these and when you are ready to start, you should click on "Click when you are ready to start" (Figure 1).

Summary of Instructions

You will be presented with a sequence of 15 different problems. In each problem you start owning some asset which you have to dispose of during or at the end of the problem. It has a value (which may be negative). When you dispose of it, the value will earn interest until the end of the problem, and this will constitute part of your payoff for that problem. In each problem there is a sequence of time periods and you can dispose of the asset in any of these. In each of the time periods until you dispose of the asset you will also earn a cash flow (which could be positive or negative) and which will be added to (or subtracted from if it is negative) your payoff from the disposal. The cash flow follows a random path, determined by Nature, jumping either up or down by a fixed specified amount each period with specified and fixed probabilities. The disposal decision can be taken in any one of the periods of the problem, though you will have to dispose of it in the final period if you have not disposed of it before then. More specifically, your task is to decide in each time period in each of the problems whether you want to continue holding the asset to the next period of the problem, or whether you want to stop and dispose of the asset in that period. You will be given the following information in each problem: the value of the asset on disposal, the initial cash flow, the size of the jump in the cash flow, the probability that the cash flow jumps up (and the residual probability that the cash flows move down), the number of periods in the problem at the end of which you have to dispose of it, and the rate of interest on the disposal value.

Cash Flows and Probabilities

If you dispose of the asset now the value of the disposal will be 909 points. The cash flows that you have gained until now are equal to 10 points.

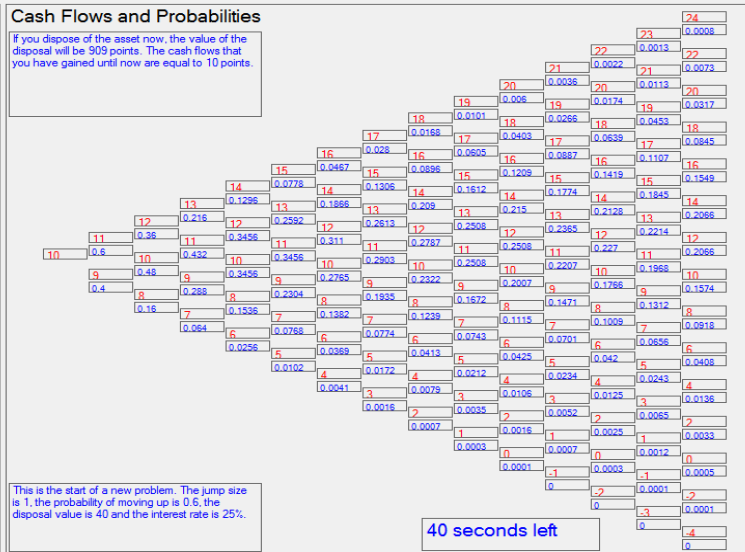


Figure 2

When you click to start the experiment, a picture like the Figure 2 appears. On the left of the screen is a summary of the instructions, while on the right appears the decision tree. The red numbers on the decision tree are the possible cash flows in the disposal problem in each of the periods of the problem, and the blue numbers are the corresponding probabilities. At the bottom under the tree are two boxes. If you are in the first period of a problem the first box tells you the jump in the cash flows, the probability of jumping up (the probability of moving down is the residual from 1) in that problem, the disposal value of the asset and the rate of interest on this value. In subsequent periods, this box shows the number of the current period, the decision by Nature as to whether the cash flow has jumped up, and the implied cash flow in this period. The second box shows the time that you have left to take the decision in that particular period. At the top left above the tree is a box. This box tells you the disposal value plus interest if you decide to dispose of it in this period, as well as the total cash flows you have accumulated up to this period.

For example, in Figure 2, you can see that the change in the cash flows in the next period is equal to +1 or -1 with probabilities 0.6 and 0.4 respectively.

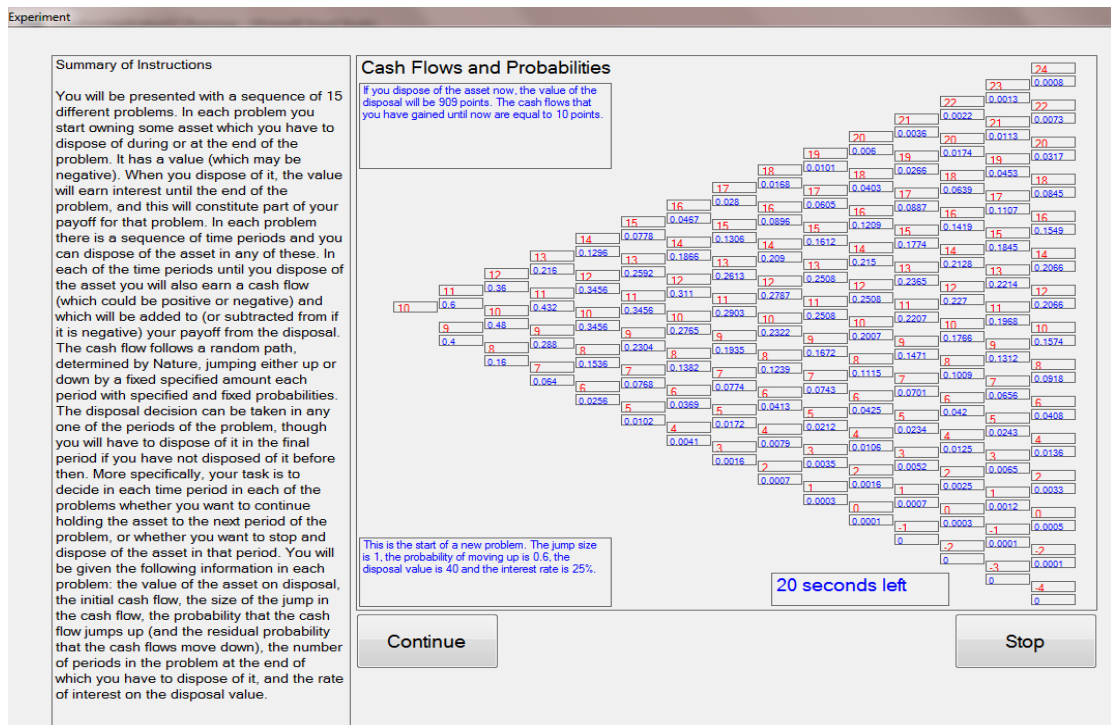


Figure 3

In each period you cannot take a decision until at least 20 seconds have elapsed; after these 20 seconds the "Continue" and "Stop" buttons appears (Figure 3). You will have a maximum time of 40 seconds in each period.

Experiment

Summary of Instructions

You will be presented with a sequence of 15 different problems. In each problem you start owning some asset which you have to dispose of during or at the end of the problem. It has a value (which may be negative). When you dispose of it, the value will earn interest until the end of the problem, and this will constitute part of your payoff for that problem. In each problem there is a sequence of time periods and you can dispose of the asset in any of these. In each of the time periods until you dispose of the asset you will also earn a cash flow (which could be positive or negative) and which will be added to (or subtracted from if it is negative) your payoff from the disposal. The cash flow follows a random path, determined by Nature, jumping either up or down by a fixed specified amount each period with specified and fixed probabilities. The disposal decision can be taken in any one of the periods of the problem, though you will have to dispose of it in the final period if you have not disposed of it before then. More specifically, your task is to decide in each time period in each of the problems whether you want to continue holding the asset to the next period of the problem, or whether you want to stop and dispose of the asset in that period. You will be given the following information in each problem: the value of the asset on disposal, the initial cash flow, the size of the jump in the cash flow, the probability that the cash flow jumps up (and the residual probability that the cash flows move down), the number of periods in the problem at the end of which you have to dispose of it, and the rate of interest on the disposal value.

Cash Flows and Probabilities

If you dispose of the asset now, the value of the disposal will be 728 points. The cash flows that you have gained until now are equal to 13 points.

Nature has chosen to move DOWN. The cash flow in this period is 9. This is period 2.

19 seconds left

Figure 4

When you have taken your decision you should click on “Continue” or “Stop” as appropriate, and then, when you are sure about your decision, you should click on “Click to confirm” as shown in Figure 4. You can change your mind about whether to stop or continue as many times as you want, as long as you have not clicked “Click to confirm” and you still have time. Notice that if you continue until the final period the disposal decision is then compulsory.

Experiment

Summary of Instructions

You will be presented with a sequence of 15 different problems. In each problem you start owning some asset which you have to dispose of during or at the end of the problem. It has a value (which may be negative). When you dispose of it, the value will earn interest until the end of the problem, and this will constitute part of your payoff for that problem. In each problem there is a sequence of time periods and you can dispose of the asset in any of these. In each of the time periods until you dispose of the asset you will also earn a cash flow (which could be positive or negative) and which will be added to (or subtracted from if it is negative) your payoff from the disposal. The cash flow follows a random path, determined by Nature, jumping either up or down by a fixed specified amount each period with specified and fixed probabilities. The disposal decision can be taken in any one of the periods of the problem, though you will have to dispose of it in the final period if you have not disposed of it before then. More specifically, your task is to decide in each time period in each of the problems whether you want to continue holding the asset to the next period of the problem, or whether you want to stop and dispose of the asset in that period. You will be given the following information in each problem: the value of the asset on disposal, the initial cash flow, the size of the jump in the cash flow, the probability that the cash flow jumps up (and the residual probability that the cash flows move down), the number of periods in the problem at the end of which you have to dispose of it, and the rate of interest on the disposal value.

Cash Flows and Probabilities

If you dispose of the asset now, the value of the disposal will be 728 points. The cash flows that you have gained until now are equal to 19 points.

Nature has chosen to move DOWN. The cash flow in this period is 9. This is period 2.

0 seconds left

Continue Stop

Click to confirm

Figure 5

Be careful!

You should answer before 40 seconds have elapsed. In the case that you run out of time a message box appears (like in Figure 5) that the time is up. In this case, the program assumes that you continue to the next period of the problem.

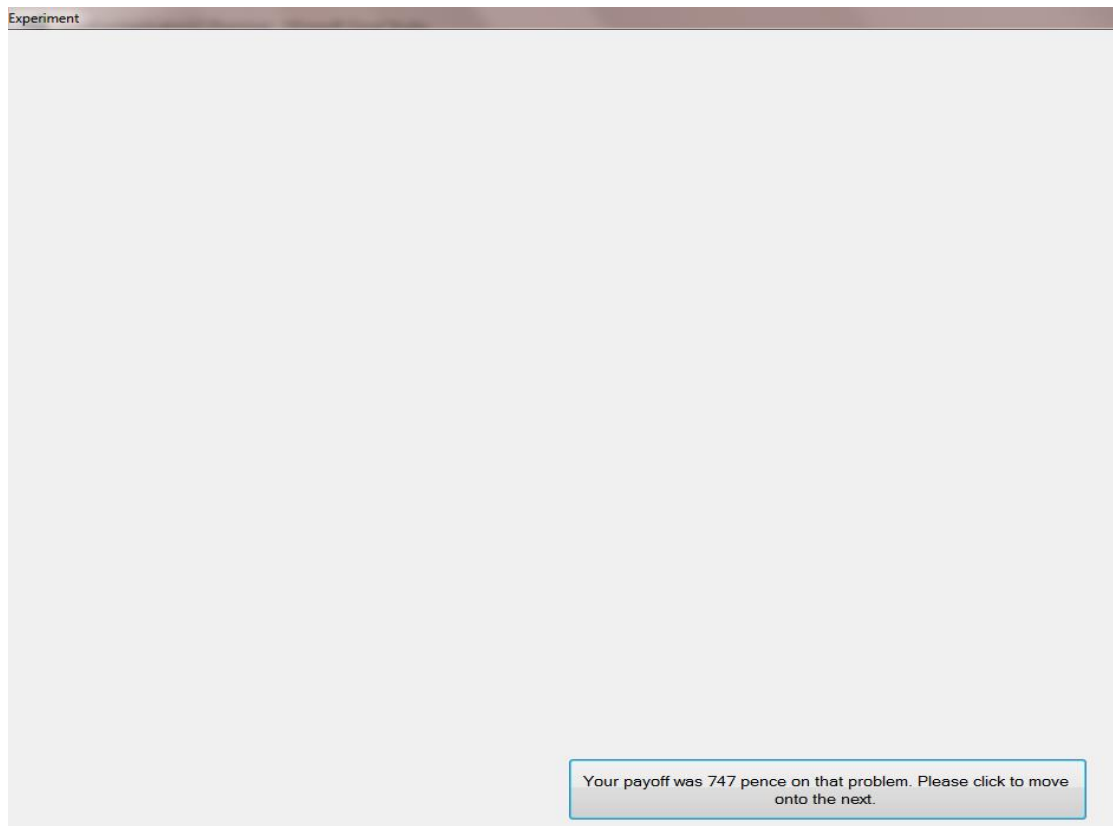


Figure 6

When you take the decision to dispose of the asset, by clicking on “Stop” and then on “Click to confirm”, a box like that in Figure 6 appears informing you about the total payment in that liquidation problem.

At the end of the experiment a message will appear asking you to call over the experimenter. Please click on this message to let us know that you have finished the experiment.

Example

By following the numbers in the Figures we are going to present an example. By looking at Figure 2, we know that the initial cash flow is equal to 10, the probability the cash flow of moving up in the next period, that is + 1, is equal to 0.6 while the probability of moving down, that is -1, is equal to 0.4. The disposal value is equal to 40 and the interest rate on it is equal

to 25%. You can also notice that the disposal value with the interest rate is equal to 909 (be careful, this value changes during the periods) and the accumulated cash flows are equal to 10, i.e. the cash flow of the first period. After the passage of the first 20 seconds, you can decide to continue or stop (Figure 3). Let us assume that you decide to continue in the next period. You just click "Continue" and "Click to confirm". Let us also assume that Nature decided the cash flow to go down in period two. Therefore, the cash flow earned in the second period is equal to 9 (Figure 4), while your accumulated cash flows are equal to 19 (= 10 + 9) and the disposal value with the interest rate is equal to 728. If you decide to Stop in the second period, your payoff is going to be equal to 747 (= 19 + 728).

How long the experiment will last

We expect you to be in the laboratory no more than one and a half hours.

Payment

Your payment from the experiment will be your payment in one randomly-chosen problem of the experiment; you will randomly choose one numbered disk from a bag containing 15 disks numbered from 1 to 15, and the number on the disk chosen will determine the problem on which you will be paid. In the experiment payments are denominated in *tokens*. These tokens will be converted into real money using the exchange rate: 100 tokens = £1. If the payment in the randomly chosen problem is negative, this will be negative.

The show up fee is £2.50 and this will be added to your payment from the experiment, described above. In no circumstances will your payment be negative.

If you have any questions, please raise your hand and an experimenter will come to you.

Chapter 3

Optimal Timing of the Exercising of a Financial Option Contract¹²

3.1 Introduction

In recent years, physicists and mathematicians have entered the fascinating world of financial derivatives. Their influence has led to many complex financial products, such as options and mortgage-backed securities. More specifically, these include some of the most complicated financial instruments, such as the call, put, barrier, compound and rainbow options, the collateralised mortgage obligations (CMOs) and the collateralised debt obligations (CDOs). Nowadays, all the above play a key role in financial markets.

We decided to experimentally explore one of the most commonly-used of these instruments: the American call option. More specifically, we wanted to investigate a key component of modern finance theory: the optimal timing of the exercising of an American call option. It is a key component since it is built into option pricing theory, and that, in turn, is built into general theories of financial markets. If the theory on the optimal exercising of a call option is shown to be empirically invalid, then the whole building constructed on top of it may come crashing down.

An American call option is a financial contract which gives the holder the *right but not the obligation* to buy an asset at a pre-specified price, which is called the strike or exercise price.

¹² This chapter is based on joint work with John Hey.

The holder of this contract can exercise it, if he or she wants, any time up to the end time point, which is stated in the contract, or he or she may decide to leave the contract to expire without exercising it. His profit is equal to the price that the asset had at the exercising point minus the exercise price. The question then is when the holder of the contract should exercise it, that is, at what asset price; or equivalently, when is the *optimal time point* to stop holding the contract unexercised. The theory of the optimal stopping time answers this.

The theory can be applied to both real (tangible) and financial (intangible) assets. For the former ones, we refer to the theory as irreversible decision making under uncertainty, or real option pricing, and for the latter ones as financial option pricing.

Henry (1974), McDonald and Siegel (1986), and Dixit and Pindyck (1994) were among the first researchers who studied the optimal timing of an investment in an irreversible project by using option pricing theory. A key issue is whether people indeed follow this theory, and invest “optimally” according to it. Even if the theory has been generalised and applied in many different contexts and kind of investments, it seems that empirical studies of the theory and of its extensions are rather few.

Furthermore, it seems that researchers do not study quite often experimentally this kind of theory in continuous time. In the literature, there is only one paper which tests experimentally the theory of an American call option in continuous time – that by Oprea *et al* (2009). They gave their subjects a set of problems in each of which the subject owned a call option and had to decide when to exercise it. Oprea *et al* tested the theory under the usual neoclassical finance assumption – that the decision-makers are risk-neutral. They found that the theory did not fit the experimental evidence particularly well. However, their study focused on the *learning process* of the subjects in exercising an American call option contract at the optimal exercising point. Based on their results it seems that subjects exercise their option contract closer to the optimal trigger as their experience increases.

We also work with a continuous time American call option problem. We built our software in Python and, in contrast to Oprea *et al* (2009), we use a discrete time approximation of the continuous time model to build our software and choose our parameters and our problems set, and not binomial approximations of the Brownian process. We concentrate our study on the behavioural aspect of risk aversion; that is we investigate whether risk-aversion can explain the departure from the theory (which is based on risk-neutrality as we have already noted).

More specifically, we ran a two part experiment. The first part looked at the exercising behaviour of our subjects faced with a decision as to when to exercise an option that they owned. In the second part, we elicited the risk attitude of our subjects through a set of lottery questions. We tested whether risk aversion affects the subjects' decisions. We did this in two ways. First, we used the data from the second part to elicit the risk-aversion parameter of the subjects, and tested whether this parameter explained their behaviour. Second, we estimated the risk aversion parameter which best explains the subjects' behaviour from the subject's decisions in the exercising problem.

We also collected demographic data for each of our subjects. We use this in our data analysis to test whether and how demographics play a role in this story.

3.2 Theory

We consider an American call option in continuous time. By holding this contract the holder has the *right but not the obligation* to buy an asset, for example a stock, at a pre-specified price. This is called the strike or exercise price. The holder can exercise the option at any time up to the contract's maturity date. If the holder exercises the option contract before the maturity date, his or her profit is be the difference between the price of the asset at the time

that he or she exercised the option and the exercise price. If the holder does not exercise the option the profit is zero.

It is usually assumed in continuous time theory that the asset price follows Geometric Brownian Motion (GBM) – a well-loved construct in finance theory. Under this assumption, as we will show, there is a well-defined trigger/price at which the option should be exercised.

In Section 3.2.1, we give an overview of the use of GBM in the asset pricing literature, as this will help us to understand better the nature of the movement of the asset's price. This is important as the asset price that our subjects were watching in their screens, and the one that we used to build our software, was derived from the discretisation of the continuous time model which includes the stochastic process of the GBM. Furthermore, as this is the base, this will help us to understand the main theory in Section 3.2.2 which gives us the formula for the optimal trigger, that is, the price at which the holder of the option should exercise it (according to the optimal stopping theory). The main theory can be found in Dixit and Pindyck (1994) and Peskir and Shiryaev (2006). We follow mainly Dixit and Pindyck (1994).

3.2.1 The significance of geometric Brownian motion in asset pricing

GBM is often used in continuous time pricing models. We use the following model to predict the asset's future price:

$$S_t = S_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma \varepsilon_t \sqrt{t}} \quad (1)$$

Here S_0 denotes the current price of the asset and S_t the price at time t . This equation implies that the future price of the stock depends *only* on the current price. In other words, according to this equation, we do not need to use the past prices of a stock in order to predict its future

price. This means that the above equation has the *Markov property* which is the independence of the future from the past, given the present.

A Markov process is a stochastic model that has the Markov property. Brownian motion is a well-known Markov process. We prove equation (1) by using the GBM and the Ito's Lemma.

Geometric Brownian Motion is given by

$$dS_t = \mu S_0 dt + \sigma S_0 dW_t, \quad (2)$$

where μ is the trend of the stock price, σ is its volatility, and dW_t is the increment of a Wiener process and is equal to $\varepsilon_t \sqrt{dt}$, where dt is an infinitesimally small positive number, and ε_t is a random variable which has a unit normal distribution.

A *Wiener process* (in addition to the Markov property) has the following properties:

1) it is continuous; 2) it is non-differentiable with respect to time t (because dt is infinitesimally small) ; 3) at time 0, $W_0=0$; 4) it has independent increments¹³; 5) the increment of a Wiener process follows a normal process with mean 0 and variance dt .

We notice that the Markov property is in line with the Weak Form Hypothesis (WFH) of Market Efficiency which states that all past prices of a stock are reflected in today's price. This means that the current price of a stock involves all the past information.

It seems to be empirically true that the WFH operates in markets as a result of competition between market participants. Specifically, arbitrage opportunities do not stay for long in the market, disappearing very soon due to the actions of the investors. This means that even if some investors hold private information or are informed of a specific pattern that the stock follows (after performing technical analysis), due to their actions in the market, the other

¹³ For any n times $0 < t_1 < t_2 < t_3 \dots < t_n$, the increment random variables $W_{t_1} - W_0$, $W_{t_2} - W_{t_1}$, $W_{t_3} - W_{t_2}$, ..., $W_{t_n} - W_{t_{n-1}}$ are independent.

investors become aware of their private information very soon, so at the end the information is no longer private but public. This puts an end to these trading opportunities (Hull, 2008).

As the future prices of a stock are uncertain, we use a Markov process in order to describe them. Specifically, equation (1) is using the GBM process.

Here we derive equation (1). We will use Ito's Lemma.

Ito's lemma states that for any function $f(S_t, t)$

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial S} dS + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} dS^2 + \dots$$

Using equation (2) we get

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial S} (\mu S dt + \sigma S dW_t) + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\mu S dt + \sigma S dW_t)^2$$

and hence

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial S} (\mu S dt + \sigma S dW_t) + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\mu^2 S^2 dt^2 + 2\mu\sigma S^2 dt dW_t + \sigma^2 S^2 dW_t^2) \quad (3)$$

Now, using property 5) of a Wiener process ($dW_t = \varepsilon_t \sqrt{dt} \sim N(0, dt)$), we get

$$\text{Var}(dW_t) = E(dW_t^2) - [E(dW_t)]^2$$

and hence

$$E(dW_t^2) = \text{Var}(dW_t) + [E(dW_t)]^2$$

which implies that

$$E(dW_t^2) = dt \quad (4)$$

In equation (3), as the term dt is infinitesimally small, that is, has a value approaching zero,

$dt^{\frac{3}{2}}$ and dt^2 go to zero before dt gets to zero. This is the reason that we get zero for the terms

dt^2 and $dt dW_t = \varepsilon_t dt^{\frac{3}{2}}$.

Therefore, we get:

$$df = \left(\frac{\partial f}{\partial t} + \frac{\partial f}{\partial S} \mu S + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right) dt + \frac{\partial f}{\partial S} \sigma S dW_t \quad (5)$$

By using that $f = \ln S$, the equation (5) becomes:

$$d \ln S = \left[0 + \frac{1}{S} \mu S + \frac{1}{2} \left(-\frac{1}{S^2} \right) \sigma^2 S^2 \right] dt + \frac{1}{S} \sigma S dW_t$$

which implies that

$$\ln S_t - \ln S_0 = \left(\mu - \frac{1}{2} \sigma^2 \right) (t - 0) + \sigma dW_t$$

and hence

$$\ln S_t = \ln S_0 + \left[\left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma dW_t \right] \ln e$$

and hence

$$\ln S_t = \ln S_0 + \ln e^{\left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma dW_t}$$

and hence

$$\ln S_t = \ln \left(S_0 * e^{\left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma dW_t} \right)$$

Thus

$$S_t = S_0 * e^{\left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma dW_t}$$

Now we know that $dW_t = \varepsilon_t \sqrt{dt}$

Thus finally we get equation (1):

$$S_t = S_0 e^{\left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma \varepsilon_t \sqrt{t}}$$

Note that, in addition to what we have previously stated, this implies that, as time t increases, the variance of S_t increases.

3.2.2 Optimal timing of the exercising of a financial option contract

Let us suppose that a decision-maker holds an option contract on a financial asset, for example a stock, whose price is S which evolves stochastically based on the following Ito's drift-diffusion process:

$$dS(t)=\mu(S(t),t)dt+\sigma(S(t),t)dW(t) \quad (6)$$

Let us simplify the above model by taking the case of a specific Wiener process, the standard GBM, with $\mu(S(t),t)=\mu S$ and $\sigma(S(t),t)=\sigma S$:

$$dS=\mu Sdt+\sigma SdW \quad (7)$$

where μ and σ are constant non-negative numbers.

Equation (7) represents a stochastic differential equation (SDE) where μ is the trend of the stock price, σ is the volatility of the stock price and dW is the increment of the Wiener process. More specifically, the increment of the Wiener process is:

$$dW_t=\varepsilon_t\sqrt{dt}\sim N(0,dt) \quad (8)$$

where dt is an infinitesimally small positive number, and ε_t is a unit normal random variable.

Based on the SDE of equation (7), we will find the optimal stopping time, or more precisely, the optimal value of S at which the holder should exercise the option.

The solution is derived from the fact that the holder should exercise the option when the value of it, $F(S)$, is equal to the maximum expected present value of the profit of acquiring the financial asset.

$$F(S)=\max E[(S_T-K)e^{-\rho T}] \quad (9)$$

where E is the expectation, S_T is the value of the underlying asset at the optimal time of the investment T , K is the strike price that the owner of the contract pays on exercising the option and ρ is the discount rate. We set $\delta=\rho-\mu$.

We assume that $\mu < \rho$. Otherwise, we will not be able to find the optimal value of the asset S^* , as the asset value S will be increasing in time and thus it will never be optimal to stop.

By using the Taylor series, the Bellman equation, and the three Boundary Conditions which the investment opportunity $F(S)$ must satisfy to obtain its maximum value, we can obtain the trigger value S^* and the value of the option $F(S^*)$ at the critical value S^* . These are the steps we follow below.

We do not obtain any profits before exercising the option. However, there is a gain from holding the option and not exercising it. This gain is the investment opportunity that the option offers, that is, its expected present value. As long as S is not the optimal price at which to exercise the option, the Bellman equation is the following:

$$\rho F(S)dt = E(dF) \quad (10)$$

Therefore, the expected return of the option in a time interval dt is equal to its expected present value.

In order to describe the derivative of the option value as a SDE, we apply the Taylor Series and we get:

$$dF(S) = \frac{1}{1!} \frac{\partial F(S)}{\partial S} dS + \frac{1^2}{2!} \frac{\partial^2 F(S)}{\partial S^2} dS^2 + \frac{1^3}{3!} \frac{\partial^3 F(S)}{\partial S^3} dS^3 + \dots \quad (11)$$

By denoting derivatives with primes, we get:

$$dF = F'(S)dS + \frac{1}{2}F''(S)dS^2 + \frac{1}{6}F'''(S)dS^3 + \dots \quad (12)$$

We know that the term dt is a positive number infinitesimally close to zero. This means that all the order condition terms above the second order condition are zero. Thus equation (12) is:

$$dF = F'(S)dS + \frac{1}{2}F''(S)dS^2 \quad (13)$$

In the above equation the Ito's Lemma appears applied on the derivative of the option value F .

By substituting equation (6) in equation (13) we get:

$$dF = F'(S)(\mu Sdt + \sigma SdW) + \frac{1}{2}F''(S)(\mu Sdt + \sigma SdW)^2$$

which implies

$$dF = F'(S)(\mu Sdt + \sigma SdW) + \frac{1}{2}F''(S)(\mu^2 S^2 dt^2 + 2\mu S\sigma Sdt dW + \sigma^2 S^2 dW^2)$$

$$dF = \mu SF'(S)dt + \sigma SF'(S)dW + \frac{1}{2}F''(S)(\sigma^2 S^2 dW^2) \quad (14)$$

By using the statistical properties of the GBM, $dW \sim N(0, dt)$, we get:

$$E(dW)=0 \text{ and}$$

$$E(dW^2)=Var(dW)+E(dW)^2$$

$$E(dW^2)=dt+0=dt$$

Therefore, the expectation of equation (14) is:

$$dF = \mu SF'(S)dt + 0 + \frac{1}{2}F''(S)(\sigma^2 S^2 dt)$$

which implies

$$dF = [\mu SF'(S) + \frac{1}{2}\sigma^2 S^2 F''(S)]dt \quad (15)$$

Therefore, the Bellman equation is:

$$\rho F(S)dt = [\mu SF'(S) + \frac{1}{2}\sigma^2 S^2 F''(S)]dt$$

which implies

$$\rho F(S) = \mu SF'(S) + \frac{1}{2}\sigma^2 S^2 F''(S)$$

and hence

$$\frac{1}{2}\sigma^2S^2F''(S) + \mu SF'(S) - \rho F(S) = 0 \quad (16)$$

or

$$\frac{1}{2}\sigma^2S^2F''(S) + (\rho - \delta)SF'(S) - \rho F(S) = 0 \quad (17)$$

The investment opportunity $F(S)$ must also satisfy the following three Boundary Conditions:

$$1) F(0)=0 \quad (18)$$

2) Value-Matching Condition:

$$F(S^*)=S^*-K \quad (19)$$

3) Smooth-Pasting Condition:

$$F'(S^*)=1 \quad (20)$$

The first boundary condition simply says that when the value of the stock is zero, the value of the option is also zero. At the critical value S^* , which is also called the trigger, the value-matching condition says that by exercising the option the net return is S^*-K . This condition implies that $F(S)$ is continuous at the optimal investment price S^* . Finally, the third condition secures that $F(S)$ is smooth at the optimal investment point S^* which ensures that there is not a better point to exercise the option contract.

We want to solve equation (17) and find the value of $F(S)$ by using a possible functional form.

This functional form is not proposed by chance. Equation (17) is a Cauchy-Euler equation type. To find the solution for this kind of equations, we assume that:

$$F(S)=S^\beta \quad (21)$$

Therefore:

$$F'(S)=\beta S^{\beta-1} \quad (22)$$

and

$$F''(S)=\beta(\beta-1)S^{\beta-2} \quad (23)$$

Substituting the above values for $F(S)$, $F'(S)$ and $F''(S)$ into equation (17) and we get:

$$\frac{1}{2}\sigma^2 S^2 \beta(\beta - 1)S^{\beta-2} + (\rho - \delta)S\beta S^{\beta-1} - \rho S^\beta = 0$$

which implies

$$\frac{1}{2}\sigma^2 \beta(\beta - 1)S^\beta + (\rho - \delta)\beta S^\beta - \rho S^\beta = 0$$

and hence

$$S^\beta \left[\frac{1}{2}\sigma^2 \beta(\beta - 1) + (\rho - \delta)\beta - \rho \right] = 0 \quad (24)$$

Because S is the price of the stock, $S^\beta > 0$.

Therefore:

$$Q(\beta) = \frac{1}{2}\sigma^2 \beta(\beta - 1) + (\rho - \delta)\beta - \rho = 0 \quad (25)$$

The solution of the above formula is of the following form:

$$F(S) = A_1 S^{\beta_1} + A_2 S^{\beta_2} \quad (26)$$

where A_1 and A_2 are constants and β_1 and β_2 are the known roots of the equation (25) which is similar to equation (17).

Equation (25) is an upward parabola because the coefficient of β^2 , $\frac{1}{2}\sigma^2$, is positive. Solving for β , we see that the value of β depends upon the values of the parameters σ , ρ and δ . We note that $Q(0) = -\rho$ and $Q(1) = -\delta$. Figure 1 illustrates the above.




Figure 1: Graph related to the Derivation of the Option Value formula

Therefore, β_1 is a number greater than 1 and β_2 is a negative number. It is obviously irrational to exercise when the value of S is far from the critical value S^* . Thus, we want to ensure that when the value of S is very small and approaches zero, that the option value $F(S)$ goes also to zero (and not to infinity). This is also implied by the first Boundary Condition (equation 18)). This means that we must set to zero the coefficient of the negative root β_2 in equation (26), that is, $A_2=0$.

Therefore, according to equations (18) and (26), the form of $F(S)$ is:

$$F(S) = AS^{\beta_1} \tag{27}$$

where A is a constant and β_1 is greater than one known constant number.

We will substitute the value of the option in equation (27) into the smooth-pasting condition:

$$F'(S^*) = (AS^{\beta_1})' = 1$$

and hence

$$\beta_1 AS^{\beta_1-1} = 1$$

which implies

$$\frac{A S^* \beta_1}{S^*} = \frac{1}{\beta_1}$$

and hence

$$A S^* \beta_1 = \frac{S^*}{\beta_1} \tag{28}$$

We notice that the RHS of equation (28) is equal to the RHS of equation (19), thus by substituting, we get:

$$\frac{S^*}{\beta_1} = S^* - K$$

which implies

$$S^* = \frac{\beta_1}{\beta_1 - 1} K \tag{29}$$

We substitute in equation (19) the functions of $F(S^*)$ and S^* as found in equations (27) and (29) and we get:

$$A = \frac{\frac{\beta_1}{\beta_1 - 1} K - K}{\left(\frac{\beta_1}{\beta_1 - 1} K\right)^{\beta_1}}$$

or

$$A = \frac{(\beta_1 - 1) \beta_1^{-1}}{\beta_1^{\beta_1} K \beta_1^{-1}} \tag{30}$$

We can obtain now the optimal investment value of the underlying asset S^* and the value of the option at that time $F(S^*)$ by using equations (27), (29) and (30). As we notice, equation (29) determines a wedge of $\beta_1/(\beta_1 - 1) > 1$ between the value of the stock and the strike price of the option which means that it is not optimal to exercise when the value of the stock is just equal to the exercise price.

By substituting the values that we found previously, we can verify the equation below:

$$F(S^*) + K = S^*$$

The above equation is actually the value-matching condition. We see that *the value of the option at the critical value S^** is:

$$F(S^*) = S^* - K = \frac{\beta_1}{\beta_1 - 1} K - K = K \left(\frac{\beta_1}{\beta_1 - 1} - 1 \right) = \frac{1}{\beta_1 - 1} K \quad (31)$$

Therefore, we note that the optimal price at which to exercise S^* is equal to $K + [1/(\beta_1 - 1)]K$, which simply means that in order the investment to be optimal, the net return at the investment time should be at least $[1/(\beta_1 - 1)]K$. This appears graphically in Figure 2.




Figure 2: Theory Graph of Trigger and Exercise Price formulas

3.3 Experimental design and implementation

3.3.1 Experimental design

We ran a two part experiment. The first part looked at the exercising behaviour of our subjects faced with a decision as to when to exercise an option that they owned. In the second part, we elicited the risk attitude of our subjects through a set of lottery questions. The first part of the experiment was computerised with the code written in Python; for the second part the code was written in Visual Studio.

3.3.1.1 Reproducing a continuous time problem in a laboratory

Doing an experiment on a continuous time theory is a big challenge and few do it. There were many problems that we had to solve. First, some thoughts crossed our minds. Does continuous time really exist? Is time continuous? Can we measure the price and reproduce it continuously?

We necessarily adopt a discrete time approximation to the continuous process as we cannot program the problem continuously on a computer. Hence the screen figure *looks* continuous but it is not.

Another problem is the infinite time horizon and discounting of the theory. It is obvious that we cannot reproduce these in a laboratory experiment. Kreps (1990) notes that due to the stopping probability of an investment opportunity a discount rate naturally emerges. Therefore, instead of a discount rate, we had a random stopping time in each problem or more specifically a continuation probability equivalent to the discount rate. Appendix A provides an explanation of this. Crucially, this can be done as a random horizon problem without discounting is analytically identical to an infinite horizon problem with discounting – and is a feature regularly used by experimenters. In the case that the random stopping point occurred before the subject had exercised, the profit for that problem was zero.

3.3.1.2 Telling the subjects about geometric Brownian motion

One of the main difficulties of this project was to explain to the subjects what exactly Brownian Motion¹⁴ is. If we had gone into the technical characteristics of this process in the

¹⁴ From now and on, whenever we refer to Brownian Motion, we mean Geometric Brownian Motion (GBM).

Instructions, it would have been impossible for the majority of the subjects to understand them. So we decided to *show* them Brownian motion rather than to *describe or explain* it.

To each of our subjects we gave a set of 100 problems. Specifically, we gave 4 blocks of 25 problems. We decided to make the human beings of our experiment familiar with GBM in each block of problems by giving them several specimen problems where the subject could exercise and observe the asset's price up to the random stopping time and hence try to understand through experience some of the sequence's characteristics. Specifically, the first 5 problems in any block were *practice* problems to make the subjects feel familiar with the parameters for the Brownian motion and for determining the random stopping time.

3.3.1.3 Choice of problems sets

We chose the problems in part 1 of our experiment by conducting extensive pre-experimental simulations in Matlab. We wanted to achieve as small time intervals as possible in order to be closest to the continuous time of the theory. However, we found that we had to deal with another issue. The choice of problems with extremely tiny time intervals was not possible to be implemented in Python as the software sometimes was demanding up to 73GB of memory to run! Therefore, we started increasing gradually the time intervals up to the point that there were not such issues. We managed to keep the time intervals indeed very small.

The choice of the problems set for the first part had also other difficulties. We wanted to have a specific number of values for the trend, the volatility and the stopping probability and at the same time a specific variety of average profits in each problem. The trend and the volatility were important for the Brownian Motion and the stopping probability was responsible for the random stopping point. In addition, we wanted to make sure that the

problems in Matlab were repeated exactly the same in Python. To achieve this we used the same random seed in both softwares.

Another key factor determining the choice of problems in the first part was that we had to choose the set in a way that we could infer from the data the risk-aversion level of the subjects. Therefore, the problems sets had to be built in a way that we could infer different levels of risk-aversion by having different decisions from the subjects.

After solving all the above issues, we ended up that the participants in our experiment would have to participate in four problems with different triggers and different optimal profits: 1 low profit, 2 medium profits and 1 large profit. This was needed in order to keep a good number for the main problems and avoid biases originating from the level of the optimal profits. In order to avoid biases such as order effects in the data analysis, we gave to our subjects the 4 blocks of our problems in 24 different orders. Furthermore, we allocated to each participant in our experiment different seeds for the stochastic parameter in the Brownian Motion.

For the first part our four problems were the following:

Problem	Trend	Volatility	Starting price	Strike/Exercise Price	Time Interval	Stopping Probability	Optimal Trigger S^*
1	0.026	0.28	40	40	0.002	0.6	49.9850
2	0.026	0.44	40	40	0.002	0.4	63.6428
3	0.055	0.44	40	40	0.002	0.6	57.3905
4	0.055	0.28	40	40	0.002	0.4	56.5700

Table 1: Part 1 - Problem Set

We tried to keep a good variety among the values for the three main parameters, that is, the trend, the volatility and the stopping probability. As already noted, the starting and the exercise price are the same in all problems; we chose the same values for these two parameters for two reasons: first, we wanted to avoid having an exercise price smaller than the starting price because possibly some of the subjects could have taken the exercise decision immediately after the start of each round, as they could have profit for sure¹⁵; second, if the profit was negative at the start of each problem, that is, the exercise price was higher than the initial price, there would have been a possibility that the price would not increase enough to reach a value above the exercise price during the life of the option. For all the above reasons, we wanted to start the profit at zero in all problems in order to avoid having effects from the differences in the initial values of these two parameters.

In the second part of the experiment, we chose to elicit risk aversion with the Allocation Method. The Allocation method was established by Loomes (1991). Some years later, Andreoni and Miller (2002), and Choi *et al* (2007) revived it and used it in a social and risky choice context respectively. Allocation Method appears in finance in portfolio choice problems. Through this method, the elicitation of the risk aversion parameter of a subject is obtained by observing the subject's allocation of an amount of experimental money, that the subject has been given, to various states of the world, with a specific probability for each state and given exchange rates between experimental money and real money for each of the states. There are also other well-known methods for eliciting the risk aversion parameter of the subjects which are used widely by economists. Among the others, the most established

¹⁵ As it is stated by (Brandimarte, 2006): "The main difficulty in pricing an American option is the existence of a free boundary due to *the possibility of early exercise*. To avoid arbitrage, the option value at each point in the (S, t) space cannot be less than the intrinsic value (i.e., the immediate payoff if the option is exercised)." This means that the American call option value $F(S, t)$ should satisfy the following condition $F(S, t) \geq \max\{S(t) - K, 0\}$ and given the fact that the subjects in our experiment do not pay to obtain the option contract, it is rational that the starting stock price is equal to the strike price.

ones are Holt-Laury Price Lists, pioneered originally by Holt and Laury (2002), the Pairwise Choice questions used by Hey and Orme (1994) and the Becker-DeGroot-Marschak mechanism created by Becker et al (1964)¹⁶. In Holt and Laury Price Lists subjects are presented with an ordered list of pairwise choices and in Pairwise Choice questions they are also presented with pairwise choices but not in a list and not ordered ones. In both cases subjects have to choose one of the two choices. However, these two methods do not inform us about the strength of preference between the two choices. This can be done using the Allocation method or the Becker-DeGroot-Marschak mechanism, and thus these latter are considered more informative methods for eliciting the risk aversion level of the subjects. As the Becker-DeGroot-Marschak mechanism seems to be more complicated for subjects to understand, we chose the Allocation method. As in the first part, we also needed a carefully chosen problem set that we could use to extract different risk aversion preferences from different decisions. We chose a subset of the problem set that Zhou and Hey (2016) used. The 40 different allocation problems that we used are attached in Table T1.

3.3.2 Experimental implementation

The experiment took place in the EXEC laboratory. As already noted, the experiment was in two parts. The first part was called the *Financial Option Decision Problem* and it concerned the exercising decision. The *Allocation Decision Problem* followed and it concerned the elicitation of the risk aversion of the participants. The experiment lasted approximately 2 hours and 96¹⁷ people participated.

¹⁶ Further methods are discussed in Charness et al (2013)

¹⁷ The number 96 was not chosen by chance. We repeated 4 times the 24 different orders of the 4 main problems in order to avoid any order effects.

3.3.2.1 Part 1: Financial option decision problem

We welcomed our subjects in the experiment and we gave them time to read the Instructions for the first part and asked if any clarification was needed. When most subjects were happy with the Instructions, we started the experiment. Subjects who still had not finished reading the Instructions, could still continue reading them as it was an individual decision making experiment. The Instructions given to the subjects are attached in Appendix B. A summary of the Instructions was also available on the screen during the whole of the first part of the experiment.

The first part was built in four main blocks. In each block one of our main problems was used. More specifically, in each block we repeated the problem twenty five times in a row. Each of the four blocks in the problems set kept the same values for the parameters of the Brownian Motion, the starting price, the exercise price and the random stopping point for all the twenty five rounds. Because GBM is stochastic, even when the same problem was repeated, that is, with exactly the same values for the parameters, the actual path of the asset price was different each time due to the random seed. Thus, the subjects did not observe the same asset price for the twenty five rounds but at the same time the parameters were the same. The subjects were informed that they had four main problems and the following table was appeared in the Instructions and also close to their screen to remind them which the sample rounds were. As can be seen from the table, the first five rounds of each block were practice problems; these did not count for payment.

Block number	Practice Problems	Problems for payment
1	1-5	6-25
2	26-30	31-50
3	51-55	56-75
4	76-80	81-100

Table 2: Practice and Payment Rounds

In each problem, the subject held an American call option. He or she had the right but not the obligation to exercise it during the life of the option. Figure 3 shows a screen shot from the software.

The participant could see the summary of the Instructions before the start of the problem and also during it. He or she could also see the number of the problem that he or she was going to start, thus he or she was informed if it was a practice problem.

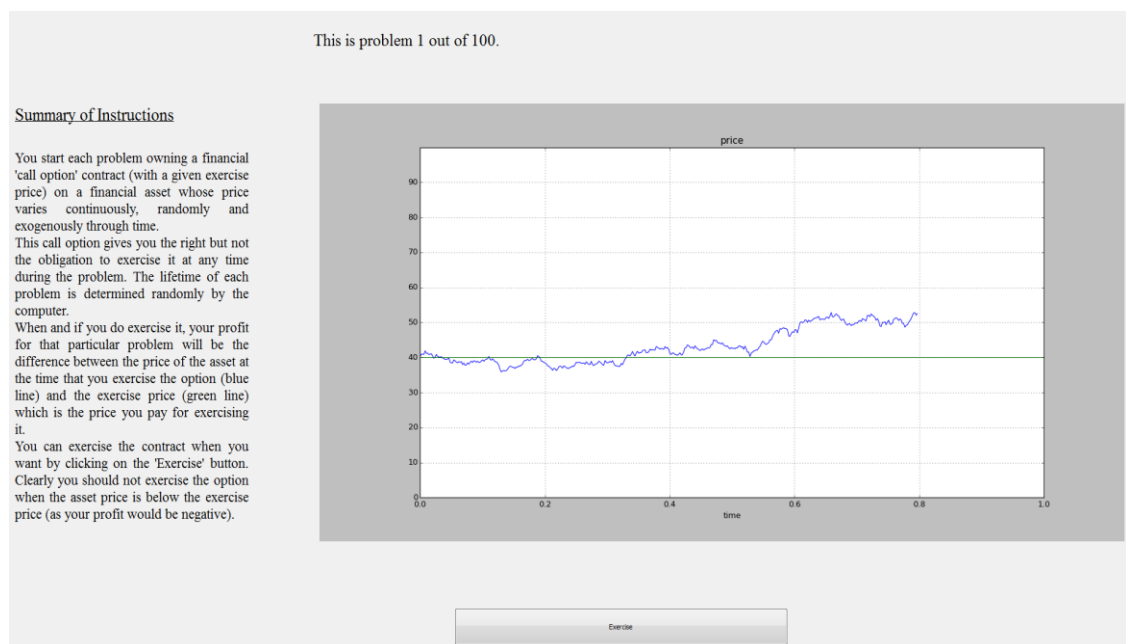


Figure 3: A Screen Shot from the First Part Software

In each problem, the subject held an option contract on a financial asset, for example, a stock, the price of which evolved continuously, randomly and exogenously.

The vertical axis represents the asset price and the horizontal one shows the time. The display adjusted in the case that the random stopping point had not occurred in the initial time interval¹⁸ or the price was over 100 units of experimental money. The green line in Figure 3 shows the exercise price and the blue one displays the asset price.

¹⁸ The time axis was actually starting adjusting after the point 0.8.

The subject had to observe the blue line and when and if he or she wanted, he or she could exercise by clicking on the “Exercise” button. After clicking, he or she could still observe the line up to the random stopping point; this was so that he or she could understand further the behaviour of the asset price.

The profit in each problem is the difference between the asset price at the exercise point and the strike price. For example, for the problem in Figure 3, if the participant had exercised slightly before point 0.8 in the time scale, he or she would have achieved as profit some 12.5 units of experimental money. We implemented the following procedure to determine the payment to a subject: he or she chose randomly one disk from a bag containing 80 disks numbered from 6 to 25, from 31 to 50, from 56 to 75 and from 81 to 100. The number on the disk chosen determined the problem on which participant was paid. Each unit of experimental money was worth £1.5. Therefore, in the above example, the profit is equal to £18.75 ($=12.5*1.5$). An example was given also to subjects. At the end of each round, the participant was informed of his profit. In the case that he or she had not exercised during the life of the option or he or she had exercised when the profit was negative¹⁹, he or she had zero profit.

3.3.2.2 Part 2: Allocation decision problem

After finishing the first part, we distributed the Instructions for the second part. As before, the subjects were given enough time to read the Instructions for this second part and they could ask questions for clarification. The screenshot from this part of the experiment is shown in Figure 4.

¹⁹ This was the case when the exercise price was higher than the asset price at the exercise time.

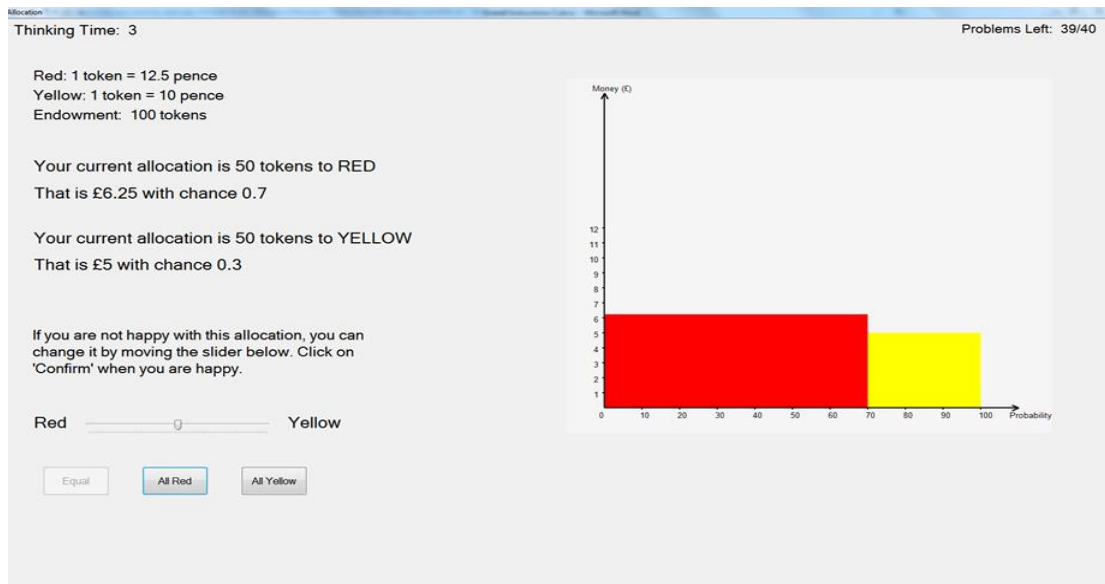


Figure 4: A Screen Shot from the Second Part Software

In this part of the experiment, in each problem the subjects had an initial endowment of 100 tokens. Their task was to allocate the tokens between two colours. Each of the colours had a specific probability of occurring.

There was an exchange rate between the allocation of tokens to each colour and the real money that the subjects would get paid depending on the outcome of the lottery; this rate was shown on the screen. The subjects were not able to take a decision until at least 5 seconds had elapsed. They could distribute the tokens among the two colours as they preferred and they could observe the exact amount of real money they would get paid given their current allocation and the exchange rates. After 5 seconds had elapsed, a “Confirm” button appeared and participants could confirm their decision when they were ready by clicking on the button.

The probability of each colour was shown on the screen. Subjects could also see this probability on the horizontal axis in the diagram and the real money they could get from their allocation and the lottery’s result on the vertical one.

As an example, let us take Figure 4. Let us assume that a subject wants to allocate equally the 100 given experimental money between the two states; 'Red' and 'Yellow'. They have respective probabilities 0.7 and 0.3. The exchange rates between tokens and money for Red is 1 token to £0.125 and for Yellow is 1 token to £0.100. The optimal allocation for a particular subject depends on the subject's attitude to risk and the exchange rates. The subject gets payment based on the outcome of a random device, for instance a lottery, in one of the problems he or she has played which is also chosen through a random device. Suppose the problem in Figure 4 was chosen among the 40 problems through a lottery. To determine the state to be realised in this problem, the subject chose one disk out of 100 disks. In the example in Figure 4, if the subject chooses any number from 1 to 70, the outcome of the allocation is Red and the participant gets $\pounds(50 \times 0.125) = \pounds 6.25$. Otherwise, a number between 71 and 100 means that the outcome of the allocation is Yellow and the corresponding payment is $\pounds(50 \times 0.1) = \pounds 5$. Overall subjects were paid the earnings on each part of the experiment plus a show up fee of £2.50.

After the end of this part a demographic questionnaire was distributed to the participants. The questionnaire is available in Appendix C.

3.4 Analysis and results

In the first part of the experiment, in each problem, whenever the participant had exercised before the expiration point, the price of the underlying asset at the exercised time was observed and saved. This price was the actual trigger of the participant for that specific problem. In all the analysis that follows we exclude from our dataset data from the practice rounds and any data from rounds in which subject's decisions to exercise would have led to negative profits (subjects knew that they would get paid zero for such a round). In Sections 3.4.1 and 3.4.2, we do not include rounds in which the subject did not exercise before the

expiration time point and thus rounds in which we were not able to observe his or her actual trigger for these rounds. In contrast, in Section 3.4.3 we do include problems in which the subject had not exercised by the time that the option expired. This latter data is informative – it tells us that the maximum price of the stock reached before the expiry of the round was lower than the subject's trigger. In Section 3.4.3 we accordingly carry out *truncated* regressions and check if the results are similar with those in Section 3.4.1.

We assumed a CARA utility function and we obtained estimates of the risk aversion parameters of each of our subjects in two ways. First, we used the data from the Allocation part of the experiment to provide an elicited risk aversion parameter for each of our subjects. Second, we obtained an estimated risk aversion coefficient of each participant through the financial option decision problems by finding the risk-aversion index which best explains the behaviour of that subject.

To estimate the best-fitting risk-aversion indices, we proceeded as follows. First, we had a program that calculated numerically (because an analytical solution was not possible) the optimal triggers for any given level of risk-aversion. Then, subject by subject, we compared the actual triggers with the optimal triggers for all levels of risk-aversion and used this to find the level of risk-aversion which best explained the subject's decisions. The first part is simple. The optimal trigger is that which maximises Expected Utility, which is simply (if we normalise the utility function so that $u(0)=0$) equal to the utility of the trigger minus the exercise price times the probability that the trigger is reached. (Note that there are just two possibilities - either the price reaches the trigger or it does not). The utility from exercising is an increasing function of the trigger while the probability that the trigger is reached is a decreasing function of the trigger. We found the latter through a simulation of the Brownian Motion. Therefore, we have not only the optimal trigger for a risk-neutral agent in each problem, but also the optimal trigger for each of the participants in each problem; first based on the

elicited risk aversion parameter and second based on the estimated risk aversion parameter. Notice that this latter definitionally fits the data best. Also, note that each of these two triggers for a specific subject is not the same in Sections 3.4.1 and 3.4.2 as no one of the two risk aversion parameters is kept the same in both Sections.

After this, we regressed the actual trigger of the subjects on the optimal triggers given both the elicited and estimated risk aversion parameters, and on some demographic data. The demographic data concern the age, the gender, the ethnicity, the education, the subject's field of study, the participant's work experience in economics/finance area, the experience in stock markets and in financial option markets, the impatience and the stress the participant felt during the experiment and the risk-averse level based on the subjects' belief. The last three factors were measured by the subjects in a scale from 1 to 5 and the level of each of these characteristics was increasing as the numbers were increasing.

3.4.1 Relation of actual subjects' trigger with risk-neutral and risk-averse optimal trigger

We began by regressing the actual trigger of the participants on the risk-neutral optimal trigger. If the theory is right we expect that the intercept should be zero and the coefficient of the risk-neutral optimal trigger should be one.

<i>Actual Trigger = $\alpha + \beta_1$ risk-neutral Optimal Trigger</i>		
	<u>Estimated Coefficient</u>	<u>Confidence Interval at 95%</u>
<i>Intercept</i>	32.8982	30.0763, 35.7202
<i>risk-neutral Optimal Trigger</i>	0.2825	0.2333, 0.3318
$R^2 = 0.0361$		

Table 3: Regression of the actual trigger on the risk-neutral optimal trigger

As can be seen, the intercept is not zero. Also, the coefficient of the risk-neutral optimal trigger is not one. Indeed the former is significantly different from zero while the latter is significantly different from one. However, this latter is positive showing that participants were increasing/decreasing their actual trigger when the one based on the theory was increasing/decreasing.

If we force the intercept to be zero, as it should be based on the theory, and run again the regression, we get Table 4:

<i>Actual Trigger = 0 + β_1 risk-neutral Optimal Trigger</i>		
	<u><i>Estimated Coefficient</i></u>	<u><i>Confidence Interval at 95%</i></u>
<i>risk-neutral Optimal Trigger</i>	0.8546	0.8502, 0.8591

Table 4: Regression of the actual trigger on the risk-neutral optimal trigger (without intercept)

As can be seen, the estimated coefficient is 0.8546 which is again positive and this time quite close to the value that the theory supports, that is one. Unfortunately the coefficient is significantly different from one.

In Table 5, a regression of the actual trigger on the *elicited* risk-averse optimal trigger (elicited from the allocation questions) shows that the coefficient of the independent variable is significantly different from zero. Unfortunately, it is also significantly different from one and the explanatory power of this variable is low. The intercept is also significantly different from zero, which it should not be.

<i>Actual Trigger = α + β_1 Elicited risk-averse Optimal Trigger</i>		
	<u><i>Estimated Coefficient</i></u>	<u><i>Confidence Interval at 95%</i></u>
<i>Intercept</i>	46.7091	44.9766, 48.4416
<i>Elicited risk-averse Optimal Trigger</i>	0.0530	0.0138, 0.0922
$R^2 = 0.0021$		

Table 5: Regression of the actual trigger on the elicited risk-neutral optimal trigger

In contrast, in Table 6, the *estimated* risk-averse optimal trigger can be seen to do much better. It seems that it influences the trigger of the subjects to an important degree. It appears that when the estimated risk-averse optimal trigger is increased by one unit the chosen trigger by the subjects is increased by 0.5278. The coefficient is significantly positive, but unfortunately it is also significantly different from 1.

<i>Actual Trigger = $\alpha + \beta_1$ Estimated risk-averse Optimal Trigger</i>		
	<u><i>Estimated Coefficient</i></u>	<u><i>Confidence Interval at 95%</i></u>
<i>Intercept</i>	25.2917	22.4236, 28.1599
<i>Estimated risk-averse Optimal Trigger</i>	0.5278	0.4642, 0.5913
$R^2 = 0.0728$		

Table 6: Regression of the actual trigger on the estimated risk-averse optimal trigger

A graphical representation of the above regressions follows below.




Figure 5: Scatter of Actual Trigger against Risk Neutral Optimal Trigger




Figure 6: Scatter of Actual Trigger against Elicited Risk Averse Optimal Trigger

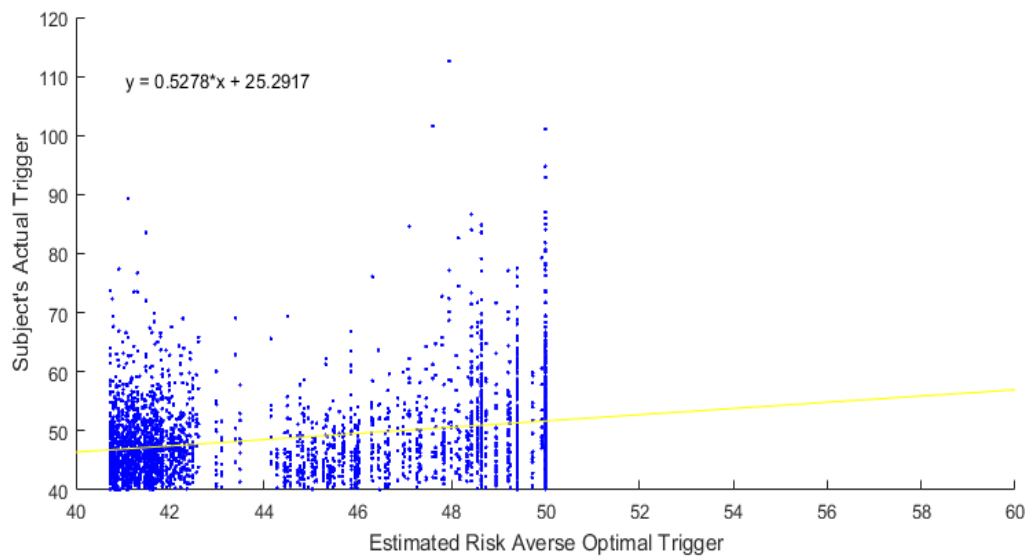


Figure 7: Scatter of Actual Trigger against Estimated Risk Averse Optimal Trigger

Based on the above results, it seems that the estimated risk-averse optimal trigger can explain behaviour better than the risk-neutral optimal trigger and the elicited risk-averse optimal one. The last comes third in the ranking as the risk-neutral optimal trigger does better. All of them have a positive relationship with the actual trigger of the participants in the experiment, showing that subjects have intuition about when the trigger should increase

or decrease, but the magnitude of the response is too low. However, it will be seen from the above graphs that subjects tend to exercise at too high a price when the optimal trigger is low, and exercise at too low a price when the optimal trigger is high: the fitted line has a positive intercept and a positive slope less than one.

3.4.2 Relation of actual subjects' trigger with risk aversion and other variables

In this Section we report a number of regressions. In each of these regressions, we use as independent variables the estimated or the elicited risk-averse optimal trigger (which is part of our main study) and one of the demographic variables (which describe demographic, educational, experience and emotional characteristics of the participants).

We first present the results for the *estimated* risk-averse optimal trigger case. An example of the kind of regressions that we have run is in Table 7.

<i>Actual Trigger = $\alpha + \beta_1$ Estimated risk-averse Optimal Trigger + β_2 Gender + β_3 Estimated risk-averse Optimal Trigger * Gender</i>		
	<i>Estimated Coefficient</i>	<i>Confidence Interval at 95%</i>
<i>Intercept</i>	24.8125	20.4246, 29.2004
<i>Estimated risk-averse Optimal Trigger</i>	0.5395	0.4434, 0.6355
<i>Gender</i>	0.9974	-4.8432, 6.8379
<i>Estimated risk-averse Optimal Trigger * Gender</i>	-0.0242	-0.1532, 0.1049
$R^2 = 0.0729$		

Table 7: Regression of the actual trigger on the estimated risk-averse optimal trigger and gender

We do the same for all the demographic variables, keeping the estimated risk-averse optimal trigger on the right hand side of the regressions. The purpose of this first series of regressions is to find the statistically significant demographic variables. We find that the statistically

significant demographic variables are: if the subjects are students; if they have studied financial options in the past; their working experience in economics or finance field; their self-perceived risk aversion level; and if they are or have been students in social sciences or natural sciences²⁰.

Then we ran again regressions like the above one in Table 7 but this time we started by including *all* of the demographic variables on the right hand side. Starting with this regression, we then dropped the least significant variable and kept the other ones, and we repeated this procedure up to the point that all of them in the regression are statistically significant. In the case with the estimated risk optimal trigger on the right hand side, we need two more regressions after the first series of regressions to end up at this point. The final regression with all the significant statistically demographic variables is presented in Table 8. It seems that the coefficient of the estimated risk-averse optimal trigger is negative but also not statistically significant from zero and not explanatory powerful in a regression with the other independent variables. First, it seems that people who have studied financial options tend to increase their trigger significantly compared to people who have not. This could possibly mean that their previous knowledge of the subject matter of the experiment make them feel more secure, that is, causing them the feeling of confidence, and thus they demand a higher trigger than the one that people with no previous educational background on financial options want. We also notice that the higher is the self-perceived risk aversion level of the subject, the lower is his actual trigger. This seems a rational result. People who study or have studied one of the natural sciences seem to be more conservative, as their actual trigger is reduced by -12.2605 in comparison with the ones who do not study these kind of sciences. One reason could be the fact that these people are more technical and they may

²⁰ Social Sciences include anthropology, ethnic and cultural studies, archaeology, area studies, economics, gender and sexuality studies, geography, organisational studies, political science, psychology and sociology. Natural Sciences include biology, chemistry, physics, earth sciences and space sciences.

understand better the characteristics of the Brownian Motion and how flexible the price of a stock can be. This thought and intuition at one point could lead them to decide to stop earlier if we take into account that the majority of people are risk-averse and thus this action could make them feel safer. We also notice that both the last two demographic variables increase the slope of the fitted line and make it again positive in total. The goodness of fit is 8.66% which may seem small but it is usual not to be really very high when we try to explain human behaviour by taking into account characteristics such as psychological ones (Frost, 2013).

<i>Actual Trigger = α + β_1 Estimated risk-averse Optimal Trigger + β_2 Study of Financial Options + β_3 Estimated risk-averse Optimal Trigger * Study of Financial Options + β_4 risk-averse Level based on Subject's Belief + β_5 Estimated risk-averse Optimal Trigger * risk-averse Level based on Subject's Belief + β_6 Natural Sciences + β_7 Estimated risk-averse Optimal Trigger * Natural Sciences</i>		
	<u>Estimated Coefficient</u>	<u>Confidence Interval at 95%</u>
Intercept	49.9334	38.8199, 61.0469
<i>Estimated risk-averse Optimal Trigger</i>	-0.0010	-0.2512, 0.2492
<i>Study of Financial Options</i>	14.7530	6.9109, 22.5951
<i>Estimated risk-averse Optimal Trigger * Study of Financial Options</i>	-0.3396	-0.5165, -0.1627
<i>risk-averse Level based on Subject's Belief</i>	-7.7649	-11.0055, -4.5243
<i>Estimated risk-averse Optimal Trigger * risk-averse Level based on Subject's Belief</i>	0.1673	0.0944, 0.2403
<i>Natural Sciences</i>	-12.2605	-21.0812, -3.4398
<i>Estimated risk-averse Optimal Trigger * Natural Sciences</i>	0.2560	0.0651, 0.4468
$R^2 = 0.0866$		

Table 8: Regression of the actual trigger on the estimated risk-averse optimal trigger and all the final statistically significant demographic variables

We repeated the same procedure this time but using the *elicited* risk-averse optimal trigger instead of the *estimated* risk-averse optimal one. By the first series of regressions find that there are more statistically significant demographic variables than there are with the estimated risk-averse optimal trigger. More specifically, we find that the statistically significant variables are: the age of the subjects: if they belong to the White category; their work experience in economics/finance; if they have studied financial options; their trading experience in stock markets; their trading experience in financial option contracts; their level of impatience during the experiment; their level of stress during the experiment; their self-perceived risk aversion level; if they are or have been students in humanities or natural sciences; and if they are or have been students in professions²¹. We follow the same procedure as in the previous case with the estimated risk-averse optimal trigger. More specifically, after running other two regressions after the first series of regressions, we end up that the statistical significant demographic variables are the ones in Table 9.

²¹ Humanities include history, linguistics, literature, performing arts, philosophy, religion and visual arts.

Professions include agriculture, architecture and design, business, divinity, education, engineering and technology, environmental studies and forestry, family and consumer science, human physical performance and recreation, journalism, media studies and communication, law, library and museum studies, medicine, military sciences, public administration, social work and transportation.

<i>Actual Trigger = α + β_1 Elicited risk-averse Optimal + β_2 Age + β_3 Elicited risk-averse Optimal Trigger * Age + β_4 White + β_5 Elicited risk-averse Optimal Trigger * White + β_6 Level of Stress + β_7 Elicited risk-averse Optimal Trigger * Level of Stress</i>		
	<u>Estimated Coefficient</u>	<u>Confidence Interval at 95%</u>
<i>Intercept</i>	41.8786	34.6817, 49.0756
<i>Elicited risk-averse Optimal Trigger</i>	0.1490	-0.0140, 0.3121
<i>Age</i>	9.4882	6.1118, 12.8647
<i>Elicited risk-averse Optimal Trigger * Age</i>	-0.2249	-0.3009, -0.1488
<i>White</i>	4.0470	0.4178, 7.6762
<i>Elicited risk-averse Optimal Trigger * White</i>	-0.0831	-0.1648, -0.0014
<i>Level of Stress</i>	-3.8980	-5.3859, -2.4101
<i>Elicited risk-averse Optimal Trigger * Level of Stress</i>	0.0968	0.0626, 0.1310
$R^2 = 0.0296$		

Table 9: Regression of the actual trigger on the elicited risk-averse optimal trigger and all the statistically significant demographic variables

We find that the coefficient of the elicited risk-averse optimal trigger is positive but not statistically significantly so. It seems that the higher is the age of the subject, the more he or she increases his actual trigger. Therefore, older people appear to be more demanding about their trigger value. People who are white seem to have a higher trigger by 4.0470 than the ones who are not. As we can notice the level of stress has negative relationship with the chosen trigger by the subject. Therefore, the more stressed is one, the lower is his actual trigger. The overall goodness of fit is very low. We finally notice that the final regression with the *estimated* risk-averse optimal trigger and the most significant demographic variables on the right hand side fits the data better than the regression in Table 9 (which uses the *elicited* risk-averse optimal trigger).

3.4.3 Relation of actual subjects' trigger with risk-neutral and risk-averse optimal trigger including the data for the rounds in which the subject had not exercised by the time the option expired

We began by implementing a truncated regression of the actual trigger of the participants on the risk-neutral optimal trigger. If the theory is right the intercept should be zero and the coefficient of the risk-neutral optimal trigger should be one.

<i>Actual Trigger = $\alpha + \beta$ risk-neutral Optimal Trigger + ϵ, where $\epsilon \sim N(0, \sigma^2)$</i>		
<i>parameter</i>	<i>Estimated Coefficient</i>	<i>Confidence Interval at 95%</i>
α	38.0352 (0.0237)	37.9888, 38.0816
β	0.2249 (0.0013)	0.2224, 0.2273
σ	7.7408 (0.0213)	7.6991, 7.7825
<i>Log-Likelihood = -12384</i>		

Table 10: Truncated regression of the actual trigger on the risk-neutral optimal trigger

As can be seen from Table 10, the intercept is not zero. Also, the coefficient of the risk-neutral optimal trigger is not one. Indeed the former is significantly different from zero while the latter is significantly different from one. However, this latter is positive showing that participants were increasing/decreasing their actual trigger when the one based on the theory was increasing/decreasing. The Log-Likelihood is very low and the estimate of σ shows us that there is lot of noise.

In Table 11, a truncated regression of the actual trigger on the *elicited* risk-averse optimal trigger (elicited from the allocation questions) shows that the coefficient of the independent variable is not statistically important to the model as it is not significantly different from zero; it is also significantly lower than one. Thus, the explanatory power of this variable is low. The intercept is significantly different from zero, which it should not be. There is lot of noise and the Log-Likelihood is lower than the one in Table 10, showing us that the regression with the

elicited risk-averse optimal trigger on the right hand side has a worse fit than the one with the risk neutral optimal trigger (Table 10).

<i>Actual Trigger = $\alpha + \beta$ Elicited risk-averse Optimal Trigger + ϵ, where $\epsilon \sim N(0, \sigma^2)$</i>		
<i>parameter</i>	<i>Estimated Coefficient</i>	<i>Confidence Interval at 95%</i>
α	47.8893 (4.2285)	39.6015, 56.1772
β	0.0681 (1.6370)	-3.1406, 3.2767
σ	7.8147 (1.0033)	5.8483, 9.7810
<i>Log-Likelihood = -12418</i>		

Table 11: Truncated regression of the actual trigger on the elicited risk-averse optimal trigger

In contrast, in Table 12, the *estimated* risk-averse optimal trigger can be seen to do much better. It seems that it influences the trigger of the subjects to an important degree. It implies that when the optimal estimated risk-averse optimal trigger is increased by one unit the chosen trigger by the subjects is increased by 0.6234. The coefficient is significantly positive, and not far away from 1 as was the case in the regression with the risk neutral optimal trigger and in the regression with the elicited risk-averse optimal trigger (Tables 10 and 11 respectively). The Log-Likelihood is larger than those in the previous two tables. However, there is still lot of noise.

<i>Actual Trigger = $\alpha + \beta$ Estimated risk-averse Optimal Trigger + ϵ, where $\epsilon \sim N(0, \sigma^2)$</i>		
<i>parameter</i>	<i>Estimated Coefficient</i>	<i>Confidence Interval at 95%</i>
α	22.6560 (0.033802)	22.5898, 22.7223
β	0.6234 (0.010142)	0.6035, 0.6433
σ	7.4934 (0.012396)	7.4691, 7.5177
<i>Log-Likelihood = -12245</i>		

Table 12: Truncated regression of the actual trigger on the estimated risk-averse optimal trigger

Graphical representations of the above regressions are shown below.

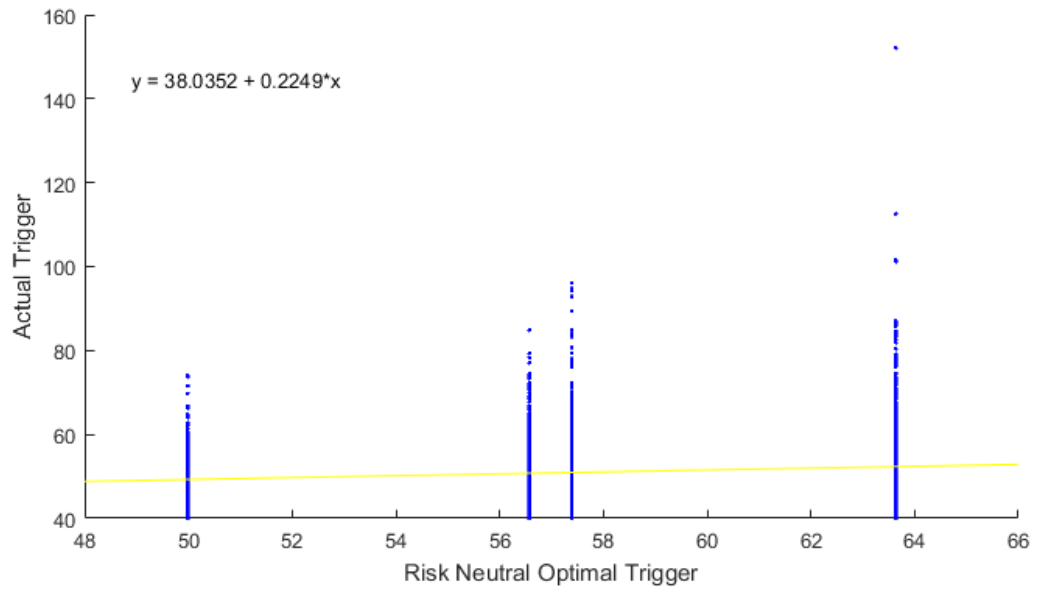


Figure 8: Scatter of Actual Trigger against Risk Neutral Optimal Trigger

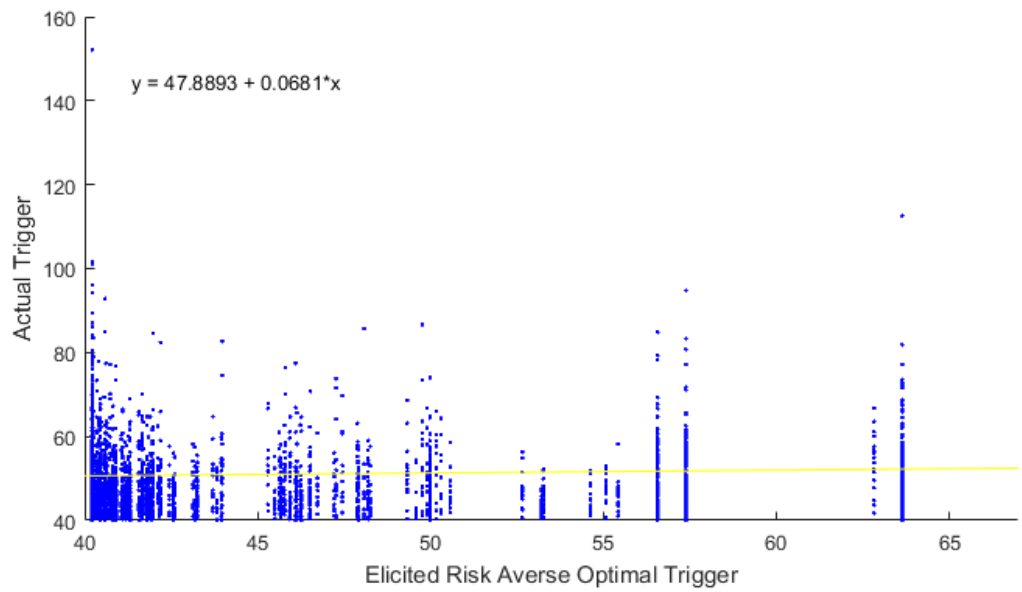


Figure 9: Scatter of Actual Trigger against Elicited Risk Averse Optimal Trigger




Figure 10: Scatter of Actual Trigger against Estimated Risk Averse Optimal Trigger

Based on the above results, it seems that the estimated risk-averse optimal trigger can explain behaviour better than the risk-neutral optimal trigger and the elicited risk-averse optimal one. The last comes third in the ranking as the regression with the risk-neutral optimal trigger fits better the data. All of them have a positive relationship with the actual trigger of the participants in the experiment, showing that subjects have intuition about when the trigger should increase or decrease, but the magnitude of the response is too low. However, it will be seen from the above graphs that the fitted line has a positive intercept and a positive slope less than one, implying that subjects tend to exercise at too high a price when the optimal trigger is low, and exercise at too low a price when the optimal trigger is high. Hence, the results confirm what we found in Section 3.4.1. concerning the ranking of the three different optimal triggers. Also, the Log-Likelihood in this Section, the R^2 in Section 3.4.1 and the values of the coefficients in both Sections lead to very similar results. Therefore, it seems that it does not make much difference whether we exclude or include the data from the rounds in which the subjects did not exercise. This is the reason that we

do not repeat the analysis of the demographic data (Section 3.4.2) using the dataset of Section 3.4.3.

3.5 Conclusions

We have tested in a laboratory experiment the theory of the timing of the exercising of an American call option. This is a crucial component of finance theory and underlies the theory of option pricing. The theory assumes that the decision-maker is risk-neutral, which we think is an implausible assumption. We therefore extended the theory to apply to risk-averse agents and tested this extension with our experimental data. Indeed we applied it in two different ways: first, by eliciting (in a separate part of the experiment) the risk-aversion of each of our subjects; and second, by finding the risk-aversion index for each subject which best explains their decisions.

The theory concerns a decision-maker who owns an American call option on some asset. This gives the decision-maker the right but not the obligation to buy one unit of the asset at a pre-specified price (called the exercise price). The profit that the decision-maker gets is the difference between the price of the asset at the time that the option is exercised and the exercise price. It is assumed in the theory that the price of the asset follows Geometric Brownian Motion (which implies that proportional changes in the price follows a random walk) in a continuous infinite-horizon world with constant discounting. The theory says that the optimal strategy of the decision-maker is to exercise the option when the price reaches some price (called the trigger price). This trigger price depends upon the mean and variance of the stochastic price process and the discount rate. We show that with risk-averse decision makers the trigger also depends upon the risk-aversion of the decision-maker – with more risk-averse agents having lower triggers, and hence exercising earlier than less risk-averse agents.

Clearly in an experiment lasting at most two hours we cannot implement an infinite-horizon constant-discounting scenario, so instead we implement an experiment with a random horizon. Specifically we had a random stopping device such that the probability of a problem terminating at the end of any period is constant. The implied continuation probability is the experimental equivalent of the discount rate. This is standard experimental practice.

We presented our subjects with 100 problems, in each of which they faced a stochastic asset price and had to decide when to exercise an option (on that asset) that they had been given. These 100 problems were 25 repetitions of 4 basic problems. The 4 basic problems had different values for the key parameters: the mean and variance of the stochastic problem and the continuation probability, chosen in such a way that there was variability in the implied optimal triggers, and the implied optimal profits.

To check if risk aversion or risk neutrality does better we used two datasets; one without including the rounds in which the subjects did not exercise and one in which we included these rounds (since we can infer that the subject's trigger was above the highest price that the stock reached during the round). The conclusions for both datasets were the same. Specifically, the first thing that we did was to see how well the risk-neutral theory described the behaviour of the subjects. The answer was 'not very well'. We then used a part of the experiment in which we had elicited the risk-aversion of our experiment to see if the elicited risk aversion explained better the behaviour of the subjects. To do this, we obviously had to calculate the implied optimal trigger for any risk-aversion. Explanation of behaviour with this elicited risk-aversion was worse. However, having knowledge of the relationship between the trigger and the level of risk-aversion enabled us to do one further thing: find, for each subject, the level of risk-aversion that best explains their behaviour. We call this level the *estimated* risk-aversion. Obviously it explains behaviour at least as well as risk-neutrality or the elicited risk-aversion. The relation between the actual trigger and the optimal trigger is

highest with the estimated risk-aversion, next highest that with risk-neutrality and lowest with the elicited risk-aversion.

We also checked the effect of some demographic variables in conjunction with the estimated or the elicited risk-averse optimal trigger on the subjects' actual triggers (we had collected such data through a questionnaire at the end of the experiment). The actual trigger regressed on the estimated risk-averse optimal trigger in conjunction with the demographic variables again does better than the regression of the actual trigger against the elicited risk-averse optimal one with the demographic variables. In the former regression, the educational background on financial options theory, the level of self-perceived risk-aversion and the study of natural sciences seem to influence their choice of the trigger.

Therefore, the risk-neutral story does not explain the behaviour of our subjects particularly well. Introducing risk-aversion into the story improves the explanatory power – though not by very much: a lot remains to be explained. There are, of course, other possible explanations. Given the context, regret and loss aversion are two obvious contenders. Subjects might suffer regret if the problem stops before they had exercised or if they had exercised early and subsequently saw the price rise. It follows that subjects could experience regret about actions they did not take and actions they took respectively. Another theory that could possibly explain why people stop earlier than the optimal time is loss aversion. Loss aversion refers to the fact that a loss counts more for someone than a profit of equal amount. Future research can also check whether regret and loss aversion explain the behaviour of our subjects better than the theories that we have tested. If so, they might be useful theories to apply in other contexts.

Tables

Table T1: Problems Set for the Allocation Decision Problem Part

Problem Number	Initial Endowment	Probability X	Probability Y	Exchange X	Exchange Y
1	100	0.1	0.9	0.57	1
2	100	0.1	0.9	0.8	1
3	100	0.1	0.9	1.25	1
4	100	0.1	0.9	1.75	1
5	100	0.2	0.8	0.5	1
6	100	0.2	0.8	0.67	1
7	100	0.2	0.8	1	1
8	100	0.2	0.8	1.5	1
9	100	0.2	0.8	2	1
10	100	0.3	0.7	0.57	1
11	100	0.3	0.7	0.8	1
12	100	0.3	0.7	1.25	1
13	100	0.3	0.7	1.75	1
14	100	0.4	0.6	0.5	1
15	100	0.4	0.6	0.67	1
16	100	0.4	0.6	1	1
17	100	0.4	0.6	1.5	1
18	100	0.4	0.6	2	1
19	100	0.5	0.5	0.57	1
20	100	0.5	0.5	0.8	1
21	100	0.5	0.5	1.25	1
22	100	0.5	0.5	1.75	1
23	100	0.6	0.4	0.5	1
24	100	0.6	0.4	0.67	1
25	100	0.6	0.4	1	1
26	100	0.6	0.4	1.5	1
27	100	0.6	0.4	2	1
28	100	0.7	0.3	0.57	1
29	100	0.7	0.3	0.8	1
30	100	0.7	0.3	1.25	1
31	100	0.7	0.3	1.75	1
32	100	0.8	0.2	0.5	1
33	100	0.8	0.2	0.67	1
34	100	0.8	0.2	1	1
35	100	0.8	0.2	1.5	1
36	100	0.8	0.2	2	1
37	100	0.9	0.1	0.57	1
38	100	0.9	0.1	0.8	1
39	100	0.9	0.1	1.25	1
40	100	0.9	0.1	1.75	1

Appendix A

Notes on connection between stopping probability and the discount factor

ρ is the *discount parameter*.

Q is the probability of expiration in a period of length 1.

We divide a time interval of 1 into $T = 1/\Delta t$ sub-intervals each of length Δt .

q is the expiration probability in an interval of length Δt .

We have

$$e^{-\rho} = 1 - Q = (1 - q)^T = (1 - q)^{1/\Delta t}$$

Taking logs we get

$$-\rho = (1/\Delta t)\ln(1 - q)$$

or

$$\rho = [-\ln(1 - q)]/\Delta t$$

ρ is used in the calculation of the optimal trigger.

The relationship between q and Q is given by $(1 - Q) = (1 - q)^T$

$$\text{So } (1 - q) = (1 - Q)^{1/T}$$

$$\text{From this it follows that } q = 1 - (1 - Q)^{1/T} = 1 - (1 - Q)^{\Delta t}$$

We note that $\rho = -\ln(1 - Q)$. Hence, the discount rate is directly tied to Q – the probability of stopping in a period of length 1. Therefore, changing dt changes q but does not change ρ .

What are the expected number of mini-periods before expiry?

event	Probability	Number of periods extant
expiration after period 1	Q	1
expiration after period 2	$(1-q)q$	2
expiration after period 3	$(1-q)^2q$	3
:		:
expiration after period n	$(1-q)^{n-1}q$	n

Check sum of probabilities = $q[1+(1-q)+(1-q)^2+\dots+(1-q)^{n-1}+\dots] = q[1/(1-1+q)] = 1$.

Expected expiry period = $q[1 + 2(1-q) + 3(1-q)^2 + \dots + n(1-q)^{n-1} + \dots$

$$= -qd[(1-q)+(1-q)^2+(1-q)^3+\dots+(1-q)^n+\dots]/dq$$

$$= -qd[\{1+(1-q)+(1-q)^2+(1-q)^3+\dots+(1-q)^n+\dots\}-1]/dq$$

$$= -qd[1/q - 1]/dq$$

$$= q/q^2 = 1/q$$

Appendix B

Experimental Instructions

The Instructions of the Experiment which were given to the subjects follow.



Welcome to this experiment. Thank you for coming. These instructions are to help you to understand what you are being asked to do during the experiment, and how you can earn money from it. This will be paid to you in cash after you have completed the experiment. The payments described below are *in addition* to the participation fee of £2.50 for the experiment as a whole that you will be paid independently of your answers.

The experiment is in two parts. The first part is about a *financial option*; the second part consists of a set of *allocation* problems. All will be explained below. The two parts are completely independent of each other. Here there are Instructions for the first part; we will distribute Instructions for the second part after you have completed the first part. We should note that the first part is estimated to take around 1 hour; the second part will last around 15 minutes, but this partly depends on how you respond.

You will be paid for both parts, and receive in addition a show-up fee of £2.50. How payments for each part will be determined is explained below.

If you have any questions on either part of this experiment, please raise your hand and an experimenter will come to you.

Konstantina Mari
John Hey

Instructions for Part 1

Financial Option Decision Problem

You will be presented with a sequence of 100 problems. In each problem, there is a financial asset whose price varies continuously, randomly and exogenously through time. You start each problem owning what is called a financial ‘call option’ contract (with a given exercise price) on the financial asset. This call option gives you the *right but not the obligation* to exercise it at any time during the problem. When and if you do exercise it, your profit for that particular problem will be the difference between the price of the asset at the time that you exercise the option and the exercise price – that is the price you pay for exercising it. Clearly you should not exercise the option when the asset price is below the exercise price (as your profit would be negative). Thus, all you have to do in each problem is to click on a button labelled ‘Exercise’ when you want to exercise the option. Of course, you may not want to exercise it and you do not have to. There is no cost to not exercising (but your profit for that problem would be zero).

The Asset Price

As we have already noted, the asset price varies continuously, randomly and exogenously through time. It is generated by the software which is programmed so that the price follows what is called in the literature ‘Geometric Brownian Motion’, which is a special case of a stochastic process. Intuitively the price sequence in each problem is such that the proportional change from one price to the next is a normally distributed random variable with constant mean and variance (which may differ across problems). An example is shown in the second figure below.

The active life of the option

In each problem, the option remains active for a random length of time, after which it can no longer be exercised. While it is active, you will be shown the path of the asset price – as in the second figure below. When it becomes inactive, the screen as in the second figure below disappears, and if you have not exercised the option by then, a message appears to inform you that your profit for that problem is zero. The lifetime of each problem is determined randomly by the computer, in such a way that it has no memory: this means that, however long it has been active, it is equally likely to be active for any given future length of time. For example, suppose that the option is still active at time 1, then, if the probability the option is still active by time 2 is some number p , then if the option is still active at time 2, then the probability the option is still active by time 3 is the same number p .

An example

Suppose the exercise price is 10; suppose you exercise the option when the asset price is 15; then your profit for that problem would be 5. Please note that these figures are stated in Experimental Currency Units (ECU). They will be converted into real money with an exchange rate that we will describe later.

The problems

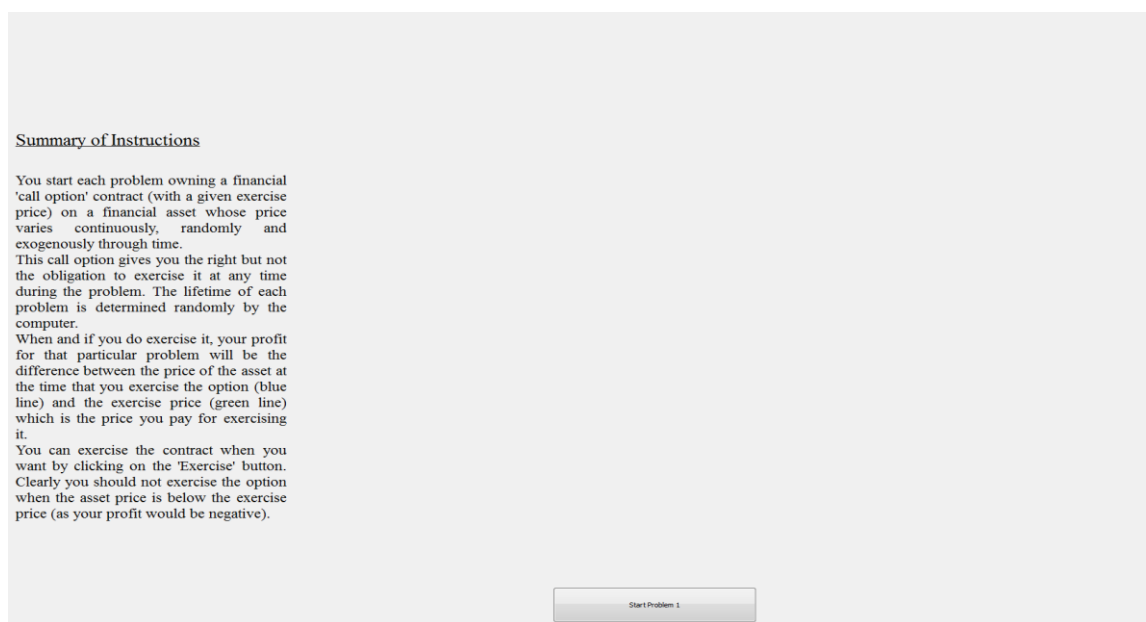
There will be 100 problems in 4 blocks each of 25. The first 5 problems in any block will be *practice problems* and will not be considered for payment; they are there to get you a feel for each problem. Each problem in any one block will have the same parameters for both the Brownian motion and for determining the random stopping time. Of course, since each

problem has both *random* motion and *random* stopping, the actual path of the asset price and the actual stopping time will be different from problem to problem. The following table shows you which problems are practice problems and which will count towards payment.

Block number	Practice Problems	Problems for payment
1	1-5	6-25
2	26-30	31-50
3	51-55	56-75
4	76-80	81-100

The Interface

When you arrive in the laboratory, you will find the screen displaying the EXEC logo. Do not touch the computer until after all the participants in the experiment have read the instructions. When all have done so, we will let you know what to do to see the first screen of the experiment, which is pictured below.



On this first screen there will be a summary of the instructions on the left of the screen and a button on the bottom right of the screen informing you which problem you are going to start. By clicking on the button, you will move on to the second screen where the problem will start. Then a screen like that below will appear:

This is problem 1 out of 100.


Summary of Instructions

You start each problem owning a financial 'call option' contract (with a given exercise price) on a financial asset whose price varies continuously, randomly and exogenously through time.

This call option gives you the right but not the obligation to exercise it at any time during the problem. The lifetime of each problem is determined randomly by the computer.

When and if you do exercise it, your profit for that particular problem will be the difference between the price of the asset at the time that you exercise the option (blue line) and the exercise price (green line) which is the price you pay for exercising it.

You can exercise the contract when you want by clicking on the 'Exercise' button. Clearly you should not exercise the option when the asset price is below the exercise price (as your profit would be negative).



Exercise

On the left of the screen is the summary of the instructions, at the top a message informs you in which out the 100 problems you are, at the bottom is the 'Exercise' button and at the middle of the screen is the graph. On the vertical axis is the price of the asset and on the horizontal axis is time. The green line shows the exercise price of the contract. The blue line shows you the asset price and its movement. You can exercise the contract when you want by clicking on the 'Exercise' button. When you click on it, you will no longer have the option to exercise in this problem (as you would have done already) but you will continue seeing the movement of the asset price up to the point that the problem will reach the random stopping time. When the option becomes inactive, a screen will appear which will inform you also for your payoff in the problem which just finished.

How long the experiment will last

We expect this part of the experiment to last around one hour. When you have completed all 100 problems, the EXEC logo and a message will appear informing you that the first part

of the experiment is over. At this point, do not touch the computer but alert one of the experimenters. When you have finished Part 1, an experimenter will start Part 2 for you.

Payment for Part 1

You will randomly choose one numbered disk from a bag containing 100 disks numbered from 6 to 25, from 31 to 50, from 56 to 75 and from 81 to 100. The number on the disk chosen will determine the problem on which you will be paid. In the experiment profits are denominated in Experimental Currency Units (ECU). Each ECU is worth £1.50; that is 10 ECU are equivalent to £15. If the profit in the randomly chosen problem is zero or negative, this payment will be zero.

Instructions for Part 2

Allocation Decision Problem

In this part of the experiment you will be given 40 different problems. In each problem, you hold 100 tokens (the conversion between tokens and money will be displayed at the top left of your screen) which you have to decide how to allocate between two risky states. These states are represented by colours (red and yellow) and each of them has a specific chance (that is, probability) of occurring. When one of the problems is played out at the end of this Part, one of the two risky states will occur (as we will explain shortly).

The Interface of the Experiment

When an experimenter starts Part 2 of the Experiment for you, you should then click on two 'Start' buttons in sequence. Then a figure like that below will appear.



At the top right of the screen, you are told how many problems are left in this part of the experiment. At the top left of the screen, you are told the exchange rate between tokens and money in that particular problem (these exchange rates may vary from problem to problem). The screen also reminds you that you have 100 tokens to allocate between the two states (red and yellow). You can allocate them in any way that you want. **You can vary the allocation by moving the slider on the bottom left of the screen.** When a problem starts, the initial allocation is 50 tokens to red and 50 to yellow. The picture at the bottom right shows you the implications of your current allocation for the amount of money you would earn in either state – the *heights* of the red and yellow boxes. This will also be written at the left-hand side of the screen. In this example, with an equal allocation of the 100 tokens, you would earn £6.25 (50 times 12.5 pence) if red occurred and £5.00 (50 times 10 pence) if yellow occurred. The *widths* of the boxes indicate the chances of each state occurring – in this example, the chance of red is 70 out of 100 and the chance of yellow is 30 out of 100. If you do not like the implications of your current allocation, you can vary it by moving the slider.

In each problem, you cannot take a decision until at least five seconds have elapsed, but you can take as long as you like. You can see the timer on the top left of the screen (“Thinking Time”). After the first 5 seconds the ‘Confirm’ button appears (see below). When you are happy with your current allocation, you should click on ‘Confirm’.

The screenshot shows an experimental interface titled "Allocation". In the top right corner, it says "Problems Left: 39/40".

On the left side, the following text is displayed:

- Red: 1 token = 12.5 pence
- Yellow: 1 token = 10 pence
- Endowment: 100 tokens

Below this, the current allocation is shown:

- Your current allocation is 50 tokens to RED
- That is £6.25 with chance 0.7
- Your current allocation is 50 tokens to YELLOW
- That is £5 with chance 0.3

A further instruction reads: "If you are not happy with this allocation, you can change it by moving the slider below. Click on 'Confirm' when you are happy."

In the center, there is a slider between "Red" and "Yellow". The slider is currently positioned at the 50 mark.


Below the slider are three buttons: "Equal", "All Red", and "All Yellow". The "All Red" button is highlighted in blue.

On the right side, there is a bar chart. The vertical axis is labeled "Money (£)" and ranges from 0 to 12. The horizontal axis is labeled "Probability" and ranges from 0 to 100. The chart shows two bars: a red bar from 0 to 70 probability with a height of 6.25, and a yellow bar from 70 to 100 probability with a height of 5.

At the bottom right of the interface is a large "Confirm" button.

Playing out a particular problem

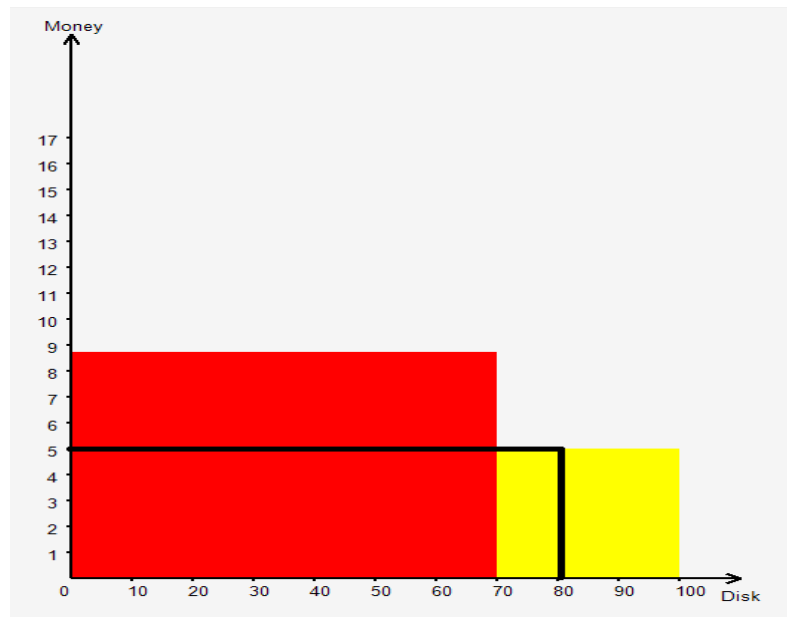
In any particular problem, you will have decided an allocation between red and yellow. This will imply a payment on this particular problem if red occurs and a payment if yellow occurs:



To determine your payment, you will draw a disk out of a bag (without looking inside the bag) containing 100 disks. The number on the disk will determine a point on the horizontal axis in the figure above; the corresponding amount on the vertical axis will be your payment from the lottery.

An Example

Let us assume you are going to play out the problem in the figure immediately above. If the number on the disk that you draw is between 1 and 70 inclusive, you earn £8.75; if the number is between 71 and 100 inclusive, you earn £5.00. It is clear that the probability of getting £8.75 is 0.7 and the probability of getting £5.00 is 0.3. Assume that the number on the disk



that you draw is 80. Then, your corresponding payment would be £5.

Payment for Part 2

You will randomly choose one numbered disk from a bag containing 40 disks numbered from 1 to 40; the number on the disk chosen will determine the problem on which lottery will take place. The experimenter will then retrieve from the computer your decision (that is, your allocation of the 100 tokens) on that problem. After this, the payout will take place as described above: you will choose a disk out of 100 disks numbered from 1 to 100 and the number on it will determine your payment.

Appendix C

Demographic questionnaire

Subject

Number:

Please provide us the following information about you.

Q. Sex: Male/Female

(Cycle the right one)

Q. Age: What is your age?

(Cycle the right one)

- Under 12 years old
- 12-17 years old
- 18-24 years old
- 25-34 years old
- 35-44 years old
- 45-54 years old
- 55-64 years old
- 65-74 years old
- 75 years or older

Q. Ethnicity origin (or Race): Please specify your ethnicity.

(Cycle the right one)

- White
- Hispanic or Latino
- Black or African American
- Native American or American Indian
- Asian / Pacific Islander
- Other²²

²² This is a combination of the other choices.

Q. Education: What is the highest degree or level of school you have completed? *If currently enrolled, highest degree received.* (Cycle the right one)

- No schooling completed
- Nursery school to 8th grade
- Some high school, no diploma
- High school graduate, diploma or the equivalent (for example: GED)
- Some college credit, no degree
- Trade/technical/vocational training
- Associate degree
- Bachelor's degree
- Master's degree
- Professional degree
- Doctorate degree

Please answer also the following questions.

Q. Are you currently a student? If so, in which level you are currently enrolled?

Q. What are you studying? Or what have you studied?

Q. Do you have any work experience in Finance or Economics? If so, for how long did/do you work in this field and which was/is your job title/titles?

Q. Have you studied financial options in the past?

Q. Have you ever traded in a stock market in reality in the past?

Q. Have you ever traded option contracts in reality in the past?

Q. Did you feel impatience during the experiment?

(Cycle the right one)

1: Not at all

2: Mainly disagree

3: Neither agree nor disagree

4: Mainly agree

5: Totally agree

Q. Did you feel stress during the experiment?

(Cycle the right one)

1: Not at all

2: Mainly disagree

3: Neither agree nor disagree

4: Mainly agree

5: Totally agree

Q. Which is your risk aversion level? From 1 to 5 the risk aversion level is increasing.

(Cycle the right one)

1 2 3 4 5

Q. What did you like in the experiment?

Q. What you did not like in the experiment?

Q. Any suggestions for improvement?

Thank you for your participation!

Chapter 4

Does the Stochastic Specification Matter?

4.1 Introduction

Underlying any statistical test of any hypothesis or any estimation of any model is some stochastic specification. It is fair to say that many economists generally pay scant attention to this, usually assuming normality somewhere. This chapter explores the implications of this, both for the hypothesis under study and the parameters being estimated.

The context in which we do this exploration is the estimation of the risk-aversion of decision-makers. This is crucial to most theories of decision-making under risk, and to many policy issues. There are several experimental methods of eliciting risk-aversion indices, the most prominent being Holt-Laury price lists (Holt and Laury 2002), pairwise choice questions (Hey and Orme 1994), the Becker-DeGroot-Marschak mechanism (Becker *et al* 1964) and allocation problems (Loomes 1991). We concentrate here on the latter method. Wilcox (2009) has done a similar analysis using the method of pairwise choice (which can be considered to be a sort of unstructured Holt-Laury price list); he concludes that the stochastic specification may well be more important than the functional specifications²³. We do not

²³ The preference functional (Expected Utility or some other), the weighting function and the utility function.

have different functional specifications, so as to concentrate on the effect of the stochastic specification.

Like all methods of eliciting preferences, one can make a variety of stochastic assumptions, but these depend on the elicitation method. Here we use allocation problems. We describe these in Section 2. In Section 3 we describe what the DM *ought* to be doing. But in experiments there is noise in subjects' behaviour. When we use experimental data to estimate their risk-aversion we need to take this noise into account. It is the description of this noise that is our stochastic specification. In Section 4 we discuss possible stochastic specifications. In order to compare between specifications we carry out extensive simulations – generating data under a variety of stochastic specifications and then estimating under them. We discuss our simulation and estimation methods in Section 5. Our results are in Section 6 and Section 7 concludes.

4.2 The allocation method for eliciting risk-aversion

This is one of a variety of methods for eliciting risk-aversion. Its advantages are that it is simple to describe and simple for subjects to understand – in contrast, for example, with the Becker-DeGroot-Marschak mechanism. One possible disadvantage, in an experimental setting, is that it implies that subjects must optimise rather than just choose, which latter is the case if the pairwise choice method is used.

In its simplest form, with the allocation method subjects are presented with a number of problems, each of the same form. In each problem, subjects are endowed with a quantity of tokens which they are asked to allocate between two *ex ante* risky states, with specified probabilities for the two states, and with specified exchange rates between tokens and real money for each of the states. Usually the experiment is computerised and the computer

records their chosen allocations on each problem. After they have responded to all the problems, one of them is chosen at random, and their allocation on that problem retrieved from the computer records. Then a random device is implemented and one of the two states is realised. The subject is paid the money value of their allocation of tokens to the realised state, using the specified exchange rate for that state.

Let us give a simple example. Suppose the two states are labelled 'Red' and 'Blue', and suppose the probabilities are respectively 0.4 and 0.6. Suppose the subject is given 100 tokens to allocate and is told that the exchange rates between tokens and money for Red is 0.8 tokens to a £1 (so that a token allocated to Red is worth £1.25) and for Blue is 1.25 tokens to a £ (so that a token allocated to Blue is worth 80p). The allocation that the subject makes is obviously dependent on his or her attitude to risk (this being the whole point of the exercise) and the exchange rates. Suppose that the subject decides to allocate 40 to Red and 60 to Blue. If this problem was randomly selected to be played out for real, and if the random device resulted in Red being selected, then the subject would be paid £(40/0.8) = £50; if the random device resulted in Blue being selected, then the subject would be paid £(60/1.25) = £48.

To estimate the level of risk-aversion – assuming that the subject obeys Expected Utility theory – a utility function should be specified and the parameter(s) of it would be estimated. To keep things simple in this simulation, we assume²⁴ a CRRA (Constant Relative Risk Aversion) utility function

$$u(x) = \frac{x^{1-r}}{1-r}$$

²⁴ An alternative would be the CARA (Constant Absolute Risk Aversion) function.

and estimate the parameter r . If this takes the value 0 the individual is risk-neutral; if positive, risk-averse and if negative, risk-loving; a higher value of r indicates a higher level of risk-aversion.

4.3 The optimal allocation

The decision made by each subject on each problem is an allocation of their endowment in that problem between the two states of the world. Let us call them State 1 and State 2. We normalise the endowment to 1, and denote the allocation to State 1 by x (so that the allocation to State 2 is $1-x$). Let us denote the probability of State 1 by p (so that the probability of State 2 is $1-p$), and the exchange rate between tokens and money in State 1 by e (normalising the exchange rate in State 2 to 1). As noted above, we assume that the DM is an Expected Utility maximiser with the CRRA utility function specified above. There is an *optimal* allocation given by the maximization of

$$EU(X) = p \frac{ex^{1-r}}{1-r} + (1-p) \frac{(1-x)^{1-r}}{1-r}$$

The solution if $r > 0$ (the DM is risk-averse) is

$$x^* = \frac{(pe)^{\frac{1}{r}}}{(pe)^{\frac{1}{r}} + (1-p)^{\frac{1}{r}} e}$$

Thus if $r > 0$ one immediate implication is that x^* is strictly bounded between 0 and 1 for non-zero p and e . Throughout this study we will use this property. This has implications for our choice of one of the stochastic specifications.

If $r \leq 0$ (the DM is risk-neutral or risk-loving) the DM will want to allocate as much as possible to one of the two states (the one depending on the problem and the exchange rate). If the

experiment puts bounds on the allocation (see below) the DM will want to allocate all or nothing.

4.4 Assumed stochastic specifications

While an optimal allocation exists for any given value of r , we assume that there is some noise in its implementation. We denote the *actual* allocation by x . We assume that the error is purely stochastic and that it has no systematic component (otherwise we would take that into account in our preference functional). It is the specification of this error that is our stochastic specification. In principle, if perhaps not in practice, the economist ought to take into account behavioural considerations when deciding on the specification. Our stochastic specifications follow.

4.4.1 Additive normal (an)

We start with the standard assumption in the economist's toolbox: that the noise is in the calculation of the optimal allocation, and that noise is added to the optimal allocation. Furthermore this noise is normal. To avoid it having a systematic component, it is assumed that this normal distribution has a zero mean. Thus there is no bias in the implementation of the allocation: on average it is equal to the optimal allocation. However it does not have a zero variance. We characterise this specification as $x = x^* + \varepsilon$ where ε is $N(0, 1/s^2)$. Following convention, we refer to s , the inverse of the standard deviation, as the *precision* of the allocation. The larger is s the lower is the magnitude of the error.

One problem with using the normal distribution is that it is unbounded. In an experiment, subjects would not be allowed to allocate more than the endowment to any one State, as

this would imply a negative allocation to the other state, and hence imply the possibility of the subject losing money. So actual decisions have to be truncated at 0 and 1.

4.4.2 Beta (*b*)

An alternative specification, and one that takes the boundedness of the optimal allocation into account, is to assume that x has a *Beta* distribution. This is bounded between 0 and 1, and has two parameters α and β . Its mean and variance are given by

$\frac{\alpha}{\alpha + \beta}$ and $\frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}$ respectively. If we put $\alpha = x^*(s-1)$ and $\beta = (1-x^*)(s-1)$, then

the mean and variance are given by x^* and $\frac{x^*(1-x^*)}{s}$ respectively. This is an attractive

specification: it means that x is unbiased, and that the noise reduces to zero as x^* approaches the bounds (0 and 1). This is behaviourally plausible: when the optimal allocation is around one-half the DM suffers the most uncertainty about it, but when it is close to 0 or 1 the DM is very sure about it. Figure 1 illustrates.

4.4.3 Random preferences normal (*rpn*)

This postulates that the noise is generated *before* the allocations are determined: the noise is in the risk-aversion parameter r . Here the story is that the DM is unsure about his or her value of r and it is this that is random. But, given any value for r , the DM implements the optimal allocation for that r without error. Notice that this is quite a different behavioural story.

Our first implementation of this random preferences story is to assume that r has a normal distribution with a given precision s . Clearly we have to specify the mean value of r (which

might be termed the DM's true value) and the precision. Note that negative values of r could be generated with this specification; these would imply risk-loving behaviour and hence all-or-nothing allocations. Again the larger is s the lower is the magnitude of the error.

4.4.4 Additive logistic (al)

This is almost identical to the additive normal, the only difference being that $x = x^* + \varepsilon$ where ε has now a *logistic* distribution with mean 0 and scale parameter $1/s$. Again we refer to s as the precision: the larger is s the lower is the magnitude of the scale of the distribution.

This distribution is very similar to the normal – but it has slightly heavier tails. Figure 2, in which the normal *pdf* is the blue dashed one, illustrates.

4.4.5 Random preferences lognormal (rpl)

This is similar to random preferences normal, though the similarities are less than between the additive normal and the additive logistic. The normal distribution is symmetrical while the log-normal is skewed to the right. Figure 3, in which the normal *pdf* is the blue dashed curve, illustrates.

4.5 The simulation and estimation program

We carried out an extensive simulation and estimation. The program, and the input files needed to run the program, can be found on the [EXEC](#) website. We ran 1000 simulations, a simulation corresponding to a decision-maker. In each simulation, we first used each of the specifications to generate some random allocations, and then we fitted under each of the

specifications. For each simulation, fitting involved the estimation of either a parameter r (being the risk-attitude for that simulation for the additive normal, beta and additive logistic specifications) or a parameter²⁵ r (being the *mean* of the random preferences normal and lognormal specifications); plus a precision parameter s .

In order to generate the random allocations we need to work with some numerically-specified preference parameters and with some specific allocation problems. The former can be found in the file *PreferenceParameters.csv*. This contains 17 lines the first 9 being a guide to the remaining 8. The data can be summarised as follows:

Parameter set	r for an, b and al	s for an	s for b	(mean) r for rpn	s for rpn	s for al	(mean) r for rpl	s for rpl
1	0.5	25	40	0.5	10	20	-0.7	5
2	0.5	50	40	0.5	20	40	-0.7	9
3	0.5	25	40	0.5	10	20	-0.7	5
4	0.5	50	40	0.5	20	40	-0.7	9
5	1.5	25	80	1.5	10	20	0.4	5
6	1.5	50	80	1.5	20	40	0.4	9
7	1.5	25	80	1.5	10	20	0.4	5
8	1.5	50	80	1.5	20	40	0.4	9

It will be seen that we chose two values for the r for an, b and al , each combined with two values for the precision s . A similar pattern was chosen elsewhere.

The simulation also requires some allocation problems. These we took from an experiment (Zhou and Hey 2016) investigating different elicitation methods. One of the methods explored was the allocation method. In that experiment, subjects were presented with 81 allocation problems (with different p 's and e 's). Here we used these 81 problems. In addition, because we were interested in the effect of the *number of problems* on the *accuracy* of the estimation, we also took a subset of 41 of these 81 problems, and we also doubled them up – thus creating a file consisting of 162 problems. In what follows, we refer to the 41 problems

²⁵ We trust that the use of the same notation will not confuse the reader.

file as the *small* data set, the 81 problems file as the medium data set and the 162 problems file as the *large* data set.

The program was written in Matlab, and used Maximum Likelihood estimation, implemented with the Matlab procedure *patternsearch*. The program produces estimates of r and s , and reports the maximised log-likelihoods.

4.6 Results

We report our results in several ways. First, we report the maximised log-likelihoods and ask the question as to whether the true specification is identified in the estimation. Then we look at the estimates of r , then the standard deviation of the estimates of r , and finally the estimates of s .

Tables 1 report the maximised log-likelihoods – our measure of goodness-of-fit: tables 1.1s through to 1.8s for the small data set; tables 1.1m through to 1.8m for the medium data set; and tables 1.1l through to 1.8l for the large data set. On each page there are 8 tables – corresponding to the eight parameter sets²⁶. The rows indicate the *true* specification – that generating the data; the columns indicate the *estimated* specification. What *should* be the case is that in each row the diagonal element should be the largest value in the row – indicating that the estimation has correctly identified the true specification. We indicate with (blue) shading the estimated specification with the highest log-likelihood in each row. It can be seen that for the small data set only 21 times out of 40 was the true specification correctly identified. Things were considerably better with the medium data set, with the true specification correctly identified 34 times out of 40; and they were marginally better again

²⁶ As there are minor differences between the results for the different parameter sets, we do not comment on them.

with the large data set – with 35 out of 40 correctly identified. Indeed, there is a greater separation of the log-likelihoods for the large data set. It should be noted that the log-likelihoods for *an* are generally very close to those for *al* – which is hardly surprising as the specifications are very close (see Figure 2). Interestingly the log-likelihoods for *b* when it is true are generally very close to those for *rpl*.

The message emerging from these tables is that, if the data set is large enough, the true specification is generally correctly identified, but that is not the case with the small data set.

Tables 2 report the *mean* value of the estimated *r* values: tables 2.1s through to 2.8s for the small data set; tables 2.1m through to 2.8m for the medium data set; and tables 2.1l through to 2.8l for the large data set. As before, on each page there are 8 tables – corresponding to the eight parameter sets. The rows indicate the *true* specification – that generating the data; the columns indicate the *estimated* specification. The column headed ‘True value’ is the *r* value used for generating the data.

Before we look at these tables we should note something about the *rpl* specification. Here the *r* value reported for *rpl* in the table above is *not* the mean of the *r* values. Here *r* is lognormally distributed with parameters *r* and *s*. This means that $\log(x)$ is normally distributed with mean *r* and precision *s*. So the mean of *r* is $\exp(r + \frac{1}{2s^2})$ and its precision is

the inverse of the square root of $\exp(\frac{1}{2s^2} - 1)\exp(2r + \frac{1}{s^2})$. In constructing tables 2, we

have taken this into account. This explains, for example, why the ‘True value’ for parameter set 1 with the small data set is 0.507, rather than the -0.7 in the table above.

The clear message that emerges from Tables 2 is that the mean *r* value is generally extremely precisely estimated. When the true *r* is 0.5, virtually all the mean estimates for the three data sets are between 0.49 and 0.51, and they are particularly close to 0.5 when the estimated specification is the true specification. When the true *r* is 1.5, virtually all the mean estimates

for the three data sets are between 1.48 and 1.52, and they are particularly close to 1.5 when the estimated specification is the true specification. In other words it seems to be the case that in this context our maximum likelihood estimates are unbiased, which is not normally necessarily the case²⁷. The message here seems to suggest that, if one is only interested in the mean value of r then the specification does not really matter.

However, we should take into account the standard deviation of the estimates of r . These are given in Tables 3: tables 3.1s through to 3.8s for the small data set; tables 3.1m through to 3.8m for the medium data set; and tables 3.1l through to 3.8l for the large data set. As before, on each page there are 8 tables – corresponding to the eight parameter sets. The rows indicate the *true* specification – that generating the data; the columns indicate the *estimated* specification. What *should* be the case is that in each row the smallest element should be along the main diagonal – where the estimated specification is the true specification. As in Tables 1 we highlight (in blue) where this is the case. It can be seen that for the small data set 29 times out of 40 this was the case; with the medium data set, 30 times out of 40; and with the large data set 32 times out of 40. Generally the standard deviations are very low (in comparison with the mean estimated r values) and importantly they decrease with the size of the problem set.

It should be remembered that the standard deviation indicates the accuracy of the estimates of r . While Tables 2 indicate that the *mean* estimates of r are very close to their true values, in practice one only has one data set (and not the 1000 in the simulation). Hence, when the standard deviation is high, individual estimates of r could depart quite significantly from the true value, but when the standard deviation is low, individual estimates of r will be generally closer to the true value. Looking at Tables 3, it is clear that the standard deviation of the r

²⁷ It can be shown that in general maximum likelihood estimates are *consistent* but not necessarily unbiased. Our results could emerge if our sample sizes are considered to be ‘close enough to infinity’.

estimates decreases with the number of problems in the data set. So size does matter: with a larger set of problems one gets more precise estimates.

If one is interested in estimates of the *precision*, a different story emerges, and here our estimates are not even unbiased – unless the estimated specification is the true specification. Tables 4 report the average value of the estimated s values: tables 4.1s through to 4.8s for the small data set; tables 4.1m through to 4.8m for the medium data set; and tables 4.1l through to 4.8l for the large data set. As before, on each page there are 8 tables – corresponding to the eight parameter sets. The rows indicate the *true* specification – that generating the data; the columns indicate the *estimated* specification. The column headed ‘True value’ is the s value used for generating the data. If one looks down the main diagonal it will be seen that when the estimated specification is the true specification, the mean s estimate is close to its true value. However, when one departs from the main diagonal significant differences emerge. So, in general, the estimated precision is quite far (and in some cases very far) from the true precision. Yet, considering Tables 2, it does not seem to be the case that misestimating s affects the mean estimates of r^{28} .

4.7 Conclusions

We should first note, and as is clear from Figures 1, 2 and 3, that our chosen specifications are *very* close. As a consequence it may not be surprising that, whatever is the estimated specification, the parameter set and the size of the problem set, the *mean* estimates of the r parameter are close to their true values. In particular it seems to be the case that the stochastic specification does not matter when it comes to estimating the *mean* level of risk-

²⁸ Incidentally this seems to be also true even if the maximum likelihood routine hits the bounds in the program – for example in Table 4.6s when rpn is the true specification and b the estimated specification.

aversion. Moreover, the effect of the size of the problem set on the *mean* of the estimates seems to be very low. However, when it comes to identifying the true specification, size does appear to have an effect: the bigger the problem set the better the identification. Size does also appear to have an effect on the standard deviation of the estimates of r : with a larger problem set one gets more precise estimates of r .

Furthermore, if one is interested in the estimates of s the precision, the specification and the size of the problem set do have a significant effect.

Does all this matter? Well, it depends on what use is to be made of the estimates: if one is going to use them for prediction of the optimal allocation with associated confidence intervals, then getting estimates correct is crucial. Figure 4 illustrates: this reports the distribution of 100,000 simulated allocations for particular preference parameters (set 1), a particular problem (number 22 from the small data set), and a particular true specification. The true specification here is b . The thick black curve is the (kernel²⁹) density function for the implied distribution using the true parameters, and the thin black curve is the density function for the implied distribution using the estimated parameters with b estimated; the green curve is that with an estimated; the blue curve with rpn estimated; the yellow curve with al estimated; and the magenta curve with rpl estimated. The complete set (consisting of 1640³⁰ graphs), can be found on the [EXEC website](#).

This Figure shows that when the estimated specification is the true specification, the distribution based on the estimated parameters is very close to that based on the true parameters, but when the estimated specification is *not* the true specification, the distribution can be quite different: when rpn is estimated the distribution is biased and

²⁹ Kernel density estimation (KDE) is a non-parametric way to estimate the probability density function of a random variable. We did not want to impose particular functional forms on the density functions, particularly as the distributions are truncated.

³⁰ Composed of the combination of 8 parameter sets, 41 problems (the small data set) and 5 true models.

skewed – leading to biased and skewed predictions; when rpl is estimated the distribution is also biased and skewed, though less so; with an and al the distributions are unbiased but have too small a spread. All the non-true distributions are quite different from the true distribution, and so will be any predictions and their associated confidence intervals.

So our conclusion must be: the stochastic specification *does* matter.

Figures

Figure 1: Beta distributions of x for $x^*=0.05$ (green on the left), 0.25 red (in the middle) and 0.5 blue (on the right); $s=50$




Figure 2: Normal (blue dashed) and logistic (red solid) distributions




Figure 3: Normal (blue dashed) and lognormal (red solid) distributions





Figure 4: Simulated allocations based on preference parameters 1 and problem number 29



Key: The true specification here is b , the parameter set number 1 and the problem number 22. The thick black curve is the (kernel) density function for the beta distribution using the true parameters, and the thin black curve is the function for the beta distribution using the estimated parameters; the green curve is that with an estimated; the blue curve with rpn estimated; the yellow curve with al estimated; the magenta curve with rpl estimated.

Table 1.1s: Log-Likelihoods parameter set 1, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-110.89	-111.56	-114.33	-111.11	-108.87
	<i>b</i>	-126.78	-115.67	-133.00	-126.03	-115.10
	<i>rpn</i>	-106.47	-101.21	-91.07	-104.69	-92.09
	<i>al</i>	-142.01	-136.71	-145.54	-141.14	-120.99
	<i>rpl</i>	-105.77	-100.06	-91.66	-104.09	-90.81

Table 1.2s: Log-Likelihoods parameter set 2, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-84.18	-89.90	-93.79	-84.30	-90.39
	<i>b</i>	-126.73	-115.87	-144.76	-125.75	-114.75
	<i>rpn</i>	-78.81	-74.40	-66.09	-76.61	-66.39
	<i>al</i>	-115.48	-115.08	-118.89	-114.75	-111.01
	<i>rpl</i>	-82.25	-78.19	-69.83	-80.67	-69.74

Table 1.3s: Log-likelihoods parameter set 3, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-110.98	-111.43	-114.56	-111.27	-109.19
	<i>b</i>	-126.84	-115.84	-130.55	-125.85	-115.06
	<i>rpn</i>	-106.32	-101.55	-90.95	-104.68	-91.69
	<i>al</i>	-141.78	-136.87	-149.03	-141.19	-121.10
	<i>rpl</i>	-105.74	-100.32	-92.13	-104.33	-91.13

Table 1.4s: Log-likelihoods parameter set 4, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-84.20	-90.21	-91.47	-84.30	-90.55
	<i>b</i>	-126.62	-115.86	-139.78	-125.85	-114.66
	<i>rpn</i>	-78.70	-73.88	-65.87	-76.90	-66.48
	<i>al</i>	-115.55	-114.66	-117.21	-114.67	-110.81
	<i>rpl</i>	-82.41	-78.32	-69.94	-80.90	-69.74

Table 1.5s: Log-Likelihoods parameter set 5, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-113.78	-114.56	-152.90	-114.64	-115.66
	<i>b</i>	-124.20	-123.84	-211.65	-124.76	-121.92
	<i>rpn</i>	-63.96	-65.85	-53.09	-62.77	-52.82
	<i>al</i>	-146.97	-149.45	-405.17	-147.02	-129.10
	<i>rpl</i>	-107.42	-109.41	-92.42	-105.61	-91.63

Table 1.6s: Log-Likelihoods parameter set 6, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-86.07	-86.34	-117.73	-86.43	-96.37
	<i>b</i>	-124.29	-123.67	-546.80	-124.63	-121.92
	<i>rpn</i>	-38.97	-50.82	-32.51	-38.47	-32.61
	<i>al</i>	-119.18	-119.54	-384.26	-118.42	-118.35
	<i>rpl</i>	-84.15	-85.47	-70.32	-82.38	-69.95

Table 1.7s: Log-Likelihoods parameter set 7, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-114.21	-114.46	-142.21	-114.68	-115.53
	<i>b</i>	-124.31	-123.60	-221.51	-124.83	-121.80
	<i>rpn</i>	-64.02	-65.98	-52.81	-62.88	-53.06
	<i>al</i>	-146.58	-149.30	-406.93	-146.58	-129.34
	<i>rpl</i>	-107.22	-109.62	-92.57	-105.69	-91.69

Table 1.8s: Log-Likelihoods parameter set 8, small data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-86.03	-86.42	-109.32	-86.22	-96.77
	<i>b</i>	-124.32	-124.03	-518.55	-124.50	-122.06
	<i>rpn</i>	-39.02	-50.74	-32.61	-38.49	-32.55
	<i>al</i>	-118.92	-120.04	-369.84	-118.42	-117.80
	<i>rpl</i>	-83.92	-85.91	-70.10	-82.32	-70.04

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 1.1m: Log-Likelihoods parameter set 1, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-220.69	-222.91	-230.90	-221.02	-224.16
	<i>b</i>	-252.02	-230.36	-282.04	-250.14	-239.71
	<i>rpn</i>	-213.20	-204.68	-181.86	-208.80	-184.11
	<i>al</i>	-282.07	-272.65	-294.85	-280.24	-246.95
	<i>rpl</i>	-211.69	-201.26	-183.42	-207.99	-181.92

Table 1.2m: Log-Likelihoods parameter set 2, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-167.56	-180.21	-183.18	-167.90	-189.31
	<i>b</i>	-251.89	-230.62	-340.76	-249.90	-241.66
	<i>rpn</i>	-157.64	-149.79	-132.88	-153.87	-133.42
	<i>al</i>	-229.73	-229.14	-247.31	-228.20	-227.37
	<i>rpl</i>	-165.52	-157.22	-140.64	-161.48	-140.13

Table 1.3m: Log-likelihoods parameter set 3, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-220.49	-222.84	-230.92	-221.00	-224.46
	<i>b</i>	-252.39	-230.88	-279.27	-250.07	-239.41
	<i>rpn</i>	-213.67	-204.69	-181.73	-209.80	-183.51
	<i>al</i>	-281.77	-273.29	-291.37	-280.41	-247.65
	<i>rpl</i>	-211.55	-201.28	-183.85	-207.96	-182.32

Table 1.4m: Log-likelihoods parameter set 4, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-167.42	-179.89	-182.95	-167.80	-188.14
	<i>b</i>	-252.53	-230.42	-337.85	-250.10	-240.57
	<i>rpn</i>	-158.25	-149.71	-132.55	-153.90	-133.18
	<i>al</i>	-229.72	-229.03	-246.87	-227.97	-227.61
	<i>rpl</i>	-165.74	-157.08	-140.66	-162.20	-140.28

Table 1.5m: Log-Likelihoods parameter set 5, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-226.20	-227.38	-290.60	-227.29	-241.29
	<i>b</i>	-246.84	-246.06	-414.46	-247.83	-251.37
	<i>rpn</i>	-127.70	-131.54	-104.82	-124.81	-104.85
	<i>al</i>	-291.93	-297.20	-746.99	-291.26	-260.27
	<i>rpl</i>	-213.82	-218.20	-184.05	-209.88	-181.84

Table 1.6m: Log-Likelihoods parameter set 6, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-170.86	-171.73	-225.68	-171.56	-209.29
	<i>b</i>	-246.81	-245.79	-887.16	-247.31	-251.00
	<i>rpn</i>	-77.66	-100.40	-63.97	-76.11	-64.09
	<i>al</i>	-236.54	-238.02	-717.05	-235.34	-244.19
	<i>rpl</i>	-167.41	-170.63	-139.84	-163.32	-139.22

Table 1.7m: Log-Likelihoods parameter set 7, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-226.40	-227.36	-299.06	-227.59	-241.88
	<i>b</i>	-247.28	-245.55	-382.28	-247.98	-251.54
	<i>rpn</i>	-127.77	-131.57	-104.81	-124.42	-104.86
	<i>al</i>	-291.30	-297.72	-758.16	-291.00	-261.37
	<i>rpl</i>	-213.71	-217.89	-183.83	-209.54	-182.06

Table 1.8m: Log-Likelihoods parameter set 8, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-170.84	-171.76	-234.97	-171.70	-209.72
	<i>b</i>	-246.42	-245.92	-952.76	-247.53	-250.88
	<i>rpn</i>	-78.02	-100.41	-63.79	-75.98	-63.93
	<i>al</i>	-236.07	-239.05	-689.76	-235.53	-245.12
	<i>rpl</i>	-167.21	-171.19	-139.67	-164.08	-139.26

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 1.1I: Log-Likelihoods parameter set 1, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-442.07	-447.18	-463.19	-442.41	-451.71
	<i>b</i>	-505.66	-461.76	-550.58	-501.42	-484.20
	<i>rpn</i>	-429.43	-411.16	-364.97	-419.24	-369.32
	<i>al</i>	-565.48	-547.53	-592.59	-561.18	-496.64
	<i>rpl</i>	-425.92	-403.88	-368.26	-417.60	-364.91

Table 1.3I: Log-likelihoods parameter set 3, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-442.03	-447.00	-466.02	-443.07	-452.59
	<i>b</i>	-505.51	-462.33	-549.22	-502.11	-483.96
	<i>rpn</i>	-429.93	-411.89	-364.54	-420.32	-368.61
	<i>al</i>	-565.68	-548.22	-577.19	-561.31	-497.29
	<i>rpl</i>	-425.66	-404.24	-368.58	-417.58	-365.24

Table 1.5I: Log-Likelihoods parameter set 5, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-453.77	-456.30	-594.31	-455.86	-486.87
	<i>b</i>	-494.67	-492.52	-809.77	-496.35	-506.74
	<i>rpn</i>	-258.05	-264.89	-210.19	-251.16	-210.39
	<i>al</i>	-584.58	-596.53	-1484.2	-583.30	-524.36
	<i>rpl</i>	-429.10	-437.85	-369.22	-421.34	-364.97

Table 1.7I: Log-Likelihoods parameter set 7, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>Rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-453.83	-455.97	-597.50	-455.74	-487.70
	<i>b</i>	-495.04	-492.59	-821.08	-496.22	-505.57
	<i>rpn</i>	-257.79	-265.00	-210.40	-250.06	-210.81
	<i>al</i>	-583.71	-597.72	-1496.1	-582.29	-525.17
	<i>rpl</i>	-429.37	-438.06	-368.35	-420.57	-364.91

Table 1.2I: Log-Likelihoods parameter set 2, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-335.95	-361.26	-369.68	-337.08	-383.33
	<i>b</i>	-505.42	-462.01	-663.29	-500.68	-483.92
	<i>rpn</i>	-317.02	-301.66	-266.32	-309.06	-267.72
	<i>al</i>	-460.55	-460.40	-499.27	-457.58	-457.97
	<i>rpl</i>	-333.17	-315.49	-282.31	-324.56	-281.50

Table 1.4I: Log-likelihoods parameter set 4, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-336.00	-361.56	-372.51	-336.77	-382.46
	<i>b</i>	-505.82	-462.40	-665.04	-501.44	-484.95
	<i>rpn</i>	-317.99	-300.97	-266.43	-309.47	-267.56
	<i>al</i>	-460.32	-459.92	-505.19	-457.64	-458.99
	<i>rpl</i>	-332.73	-315.62	-282.33	-325.16	-281.56

Table 1.6I: Log-Likelihoods parameter set 6, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-343.16	-344.12	-461.50	-344.54	-426.49
	<i>b</i>	-494.55	-492.75	-1753.0	-495.64	-505.12
	<i>rpn</i>	-156.94	-201.20	-129.59	-153.40	-128.75
	<i>al</i>	-473.84	-478.11	-1429.5	-471.95	-492.07
	<i>rpl</i>	-336.92	-343.50	-280.97	-327.99	-279.11

Table 1.8I: Log-Likelihoods parameter set 8, large data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	-342.74	-344.44	-470.41	-344.33	-428.16
	<i>b</i>	-494.25	-492.57	-1765.0	-496.30	-504.76
	<i>rpn</i>	-157.67	-201.08	-128.95	-153.44	-128.98
	<i>al</i>	-473.79	-479.22	-1441.2	-471.93	-494.67
	<i>rpl</i>	-336.94	-343.43	-280.40	-328.83	-279.47

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 2.1s: *r* estimates, parameter set 1, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.503	0.497	0.524	0.503	0.523
	<i>b</i>	0.500	0.502	0.500	0.539	0.496	0.535
	<i>rpn</i>	0.500	0.494	0.495	0.500	0.499	0.499
	<i>al</i>	0.500	0.519	0.483	0.569	0.515	0.566
	<i>rpl</i>	0.507	0.499	0.501	0.506	0.499	0.505

Table 2.2s: *r* estimates, parameter set 2, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.500	0.500	0.507	0.500	0.507
	<i>b</i>	0.500	0.501	0.500	0.539	0.497	0.535
	<i>rpn</i>	0.500	0.498	0.500	0.500	0.499	0.500
	<i>al</i>	0.500	0.504	0.496	0.526	0.503	0.526
	<i>rpl</i>	0.500	0.497	0.499	0.500	0.497	0.499

Table 2.3s: *r* estimates, parameter set 3, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.505	0.496	0.524	0.502	0.522
	<i>b</i>	0.500	0.502	0.500	0.539	0.497	0.536
	<i>rpn</i>	0.500	0.494	0.494	0.500	0.499	0.499
	<i>al</i>	0.500	0.522	0.481	0.568	0.513	0.567
	<i>rpl</i>	0.507	0.499	0.501	0.505	0.498	0.506

Table 2.4s: *r* estimates, parameter set 4, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.501	0.499	0.507	0.500	0.508
	<i>b</i>	0.500	0.500	0.500	0.540	0.496	0.533
	<i>rpn</i>	0.500	0.498	0.500	0.500	0.500	0.500
	<i>al</i>	0.500	0.504	0.495	0.527	0.503	0.527
	<i>rpl</i>	0.500	0.497	0.499	0.500	0.497	0.500

Table 2.5s: *r* estimates, parameter set 5, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.502	1.503	1.587	1.501	1.580
	<i>b</i>	1.500	1.501	1.501	1.615	1.495	1.614
	<i>rpn</i>	1.500	1.496	1.496	1.500	1.498	1.500
	<i>al</i>	1.500	1.511	1.483	1.704	1.508	1.726
	<i>rpl</i>	1.522	1.483	1.478	1.524	1.489	1.522

Table 2.6s: *r* estimates, parameter set 6, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.502	1.500	1.525	1.500	1.523
	<i>b</i>	1.500	1.503	1.502	1.587	1.500	1.616
	<i>rpn</i>	1.500	1.498	1.499	1.500	1.499	1.500
	<i>al</i>	1.500	1.501	1.495	1.585	1.497	1.597
	<i>rpl</i>	1.501	1.487	1.487	1.501	1.490	1.501

Table 2.7s: *r* estimates, parameter set 7, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.503	1.499	1.593	1.503	1.584
	<i>b</i>	1.500	1.503	1.506	1.617	1.497	1.622
	<i>rpn</i>	1.500	1.495	1.494	1.499	1.498	1.500
	<i>al</i>	1.500	1.511	1.486	1.685	1.506	1.716
	<i>rpl</i>	1.522	1.481	1.477	1.521	1.486	1.521

Table 2.8s: *r* estimates, parameter set 8, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.499	1.500	1.526	1.500	1.526
	<i>b</i>	1.500	1.503	1.504	1.601	1.495	1.618
	<i>rpn</i>	1.500	1.498	1.499	1.500	1.499	1.500
	<i>al</i>	1.500	1.504	1.495	1.570	1.504	1.596
	<i>rpl</i>	1.501	1.489	1.488	1.501	1.488	1.501

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 2.1m: *r* estimates, parameter set 1, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.502	0.496	0.517	0.502	0.519
	<i>b</i>	0.500	0.503	0.499	0.542	0.496	0.541
	<i>rpn</i>	0.500	0.495	0.495	0.500	0.498	0.500
	<i>al</i>	0.500	0.518	0.483	0.571	0.511	0.578
	<i>rpl</i>	0.507	0.500	0.501	0.506	0.499	0.506

Table 2.3m: *r* estimates, parameter set 3, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.503	0.495	0.516	0.502	0.521
	<i>b</i>	0.500	0.502	0.501	0.542	0.496	0.541
	<i>rpn</i>	0.500	0.494	0.494	0.500	0.499	0.500
	<i>al</i>	0.500	0.519	0.484	0.571	0.511	0.580
	<i>rpl</i>	0.507	0.499	0.502	0.506	0.499	0.507

Table 2.5m: *r* estimates, parameter set 5, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.503	1.498	1.618	1.501	1.623
	<i>b</i>	1.500	1.500	1.502	1.639	1.495	1.667
	<i>rpn</i>	1.500	1.496	1.495	1.500	1.498	1.499
	<i>al</i>	1.500	1.514	1.477	1.737	1.501	1.799
	<i>rpl</i>	1.522	1.479	1.473	1.521	1.488	1.520

Table 2.7m: *r* estimates, parameter set 7, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.501	1.499	1.616	1.500	1.622
	<i>b</i>	1.500	1.499	1.503	1.643	1.498	1.670
	<i>rpn</i>	1.500	1.494	1.494	1.500	1.498	1.500
	<i>al</i>	1.500	1.512	1.473	1.721	1.502	1.795
	<i>rpl</i>	1.522	1.479	1.473	1.522	1.487	1.522

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpl*: random preferences lognormal

Table 2.2m: *r* estimates, parameter set 2, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.501	0.500	0.502	0.500	0.502
	<i>b</i>	0.500	0.502	0.500	0.544	0.496	0.543
	<i>rpn</i>	0.500	0.498	0.500	0.500	0.499	0.500
	<i>al</i>	0.500	0.504	0.496	0.520	0.503	0.523
	<i>rpl</i>	0.500	0.498	0.499	0.499	0.497	0.499

Table 2.4m: *r* estimates, parameter set 4, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.501	0.500	0.502	0.500	0.503
	<i>b</i>	0.500	0.502	0.499	0.541	0.496	0.543
	<i>rpn</i>	0.500	0.498	0.500	0.500	0.499	0.500
	<i>al</i>	0.500	0.504	0.495	0.520	0.503	0.525
	<i>rpl</i>	0.500	0.497	0.499	0.499	0.497	0.500

Table 2.6m: *r* estimates, parameter set 6, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.501	1.500	1.539	1.501	1.542
	<i>b</i>	1.500	1.500	1.502	1.659	1.497	1.665
	<i>rpn</i>	1.500	1.498	1.499	1.500	1.499	1.500
	<i>al</i>	1.500	1.501	1.496	1.634	1.499	1.635
	<i>rpl</i>	1.501	1.487	1.486	1.501	1.489	1.501

Table 2.8m: *r* estimates, parameter set 8, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.500	1.501	1.538	1.500	1.539
	<i>b</i>	1.500	1.500	1.504	1.632	1.495	1.665
	<i>rpn</i>	1.500	1.498	1.499	1.500	1.498	1.500
	<i>al</i>	1.500	1.500	1.498	1.628	1.502	1.634
	<i>rpl</i>	1.501	1.489	1.486	1.502	1.489	1.501

Table 2.1l: *r* estimates, parameter set 1, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.502	0.496	0.514	0.502	0.521
	<i>b</i>	0.500	0.502	0.500	0.538	0.496	0.546
	<i>rpn</i>	0.500	0.494	0.495	0.500	0.498	0.500
	<i>al</i>	0.500	0.518	0.483	0.565	0.510	0.584
	<i>rpl</i>	0.507	0.500	0.502	0.506	0.499	0.507

Table 2.3l: *r* estimates, parameter set 3, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.503	0.496	0.515	0.502	0.523
	<i>b</i>	0.500	0.502	0.500	0.538	0.497	0.545
	<i>rpn</i>	0.500	0.494	0.495	0.500	0.499	0.500
	<i>al</i>	0.500	0.518	0.483	0.565	0.512	0.585
	<i>rpl</i>	0.507	0.499	0.502	0.506	0.499	0.507

Table 2.5l: *r* estimates, parameter set 5, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.501	1.497	1.616	1.502	1.635
	<i>b</i>	1.500	1.499	1.503	1.669	1.496	1.686
	<i>rpn</i>	1.500	1.496	1.494	1.500	1.498	1.500
	<i>al</i>	1.500	1.509	1.477	1.801	1.501	1.826
	<i>rpl</i>	1.522	1.480	1.473	1.522	1.486	1.521

Table 2.7l: *r* estimates, parameter set 7, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.501	1.498	1.622	1.500	1.638
	<i>b</i>	1.500	1.501	1.501	1.660	1.496	1.684
	<i>rpn</i>	1.500	1.495	1.494	1.500	1.498	1.500
	<i>al</i>	1.500	1.507	1.473	1.787	1.500	1.825
	<i>rpl</i>	1.522	1.479	1.472	1.523	1.487	1.522

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 2.2l: *r* estimates, parameter set 2, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.501	0.500	0.502	0.501	0.503
	<i>b</i>	0.500	0.502	0.500	0.546	0.497	0.545
	<i>rpn</i>	0.500	0.498	0.500	0.500	0.500	0.500
	<i>al</i>	0.500	0.504	0.495	0.518	0.503	0.526
	<i>rpl</i>	0.500	0.497	0.499	0.499	0.497	0.499

Table 2.4l: *r* estimates, parameter set 4, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.500	0.501	0.500	0.502	0.500	0.504
	<i>b</i>	0.500	0.501	0.500	0.545	0.497	0.547
	<i>rpn</i>	0.500	0.498	0.499	0.500	0.500	0.500
	<i>al</i>	0.500	0.504	0.495	0.519	0.503	0.527
	<i>rpl</i>	0.500	0.497	0.499	0.500	0.497	0.500

Table 2.6l: *r* estimates, parameter set 6, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.500	1.500	1.538	1.501	1.550
	<i>b</i>	1.500	1.499	1.502	1.734	1.497	1.685
	<i>rpn</i>	1.500	1.498	1.499	1.500	1.499	1.500
	<i>al</i>	1.500	1.502	1.496	1.681	1.500	1.648
	<i>rpl</i>	1.501	1.488	1.486	1.501	1.489	1.501

Table 2.8l: *r* estimates, parameter set 8, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	1.500	1.501	1.500	1.539	1.501	1.550
	<i>b</i>	1.500	1.502	1.501	1.744	1.496	1.679
	<i>rpn</i>	1.500	1.498	1.499	1.500	1.499	1.500
	<i>al</i>	1.500	1.499	1.497	1.665	1.502	1.651
	<i>rpl</i>	1.501	1.488	1.486	1.501	1.489	1.501

Table 3.1s: Standard deviation of r estimates, parameter set 1, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.019	0.023	0.035	0.019	0.055
	b	0.028	0.024	0.045	0.027	0.065
	rpn	0.023	0.021	0.017	0.021	0.033
	al	0.041	0.050	0.074	0.041	0.103
	rpl	0.023	0.021	0.017	0.022	0.031

Table 3.2s: Standard deviation of r estimates, parameter set 2, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.009	0.013	0.016	0.009	0.029
	b	0.029	0.025	0.046	0.028	0.064
	rpn	0.011	0.010	0.008	0.010	0.017
	al	0.021	0.026	0.036	0.020	0.062
	rpl	0.012	0.011	0.009	0.012	0.018

Table 3.3s: Standard deviation of r estimates, parameter set 3, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.018	0.024	0.033	0.019	0.055
	b	0.028	0.025	0.044	0.027	0.065
	rpn	0.022	0.021	0.016	0.022	0.034
	al	0.043	0.049	0.068	0.041	0.104
	rpl	0.022	0.020	0.017	0.021	0.034

Table 3.4s: Standard deviation of r estimates, parameter set 4, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.009	0.012	0.016	0.009	0.031
	b	0.029	0.025	0.044	0.027	0.063
	rpn	0.011	0.010	0.008	0.010	0.017
	al	0.021	0.025	0.036	0.020	0.060
	rpl	0.012	0.011	0.009	0.012	0.018

Table 3.5s: Standard deviation of r estimates, parameter set 5, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.054	0.056	0.118	0.054	0.062
	b	0.067	0.064	0.153	0.071	0.075
	rpn	0.023	0.023	0.018	0.022	0.011
	al	0.119	0.133	0.317	0.117	0.110
	rpl	0.064	0.069	0.051	0.063	0.031

Table 3.6s: Standard deviation of r estimates, parameter set 6, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.027	0.027	0.056	0.028	0.035
	b	0.068	0.064	0.221	0.066	0.078
	rpn	0.011	0.012	0.010	0.011	0.006
	al	0.062	0.065	0.207	0.055	0.066
	rpl	0.038	0.038	0.027	0.036	0.019

Table 3.7s: Standard deviation of r estimates, parameter set 7, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.055	0.053	0.125	0.056	0.061
	b	0.065	0.064	0.150	0.069	0.077
	rpn	0.022	0.024	0.017	0.021	0.012
	al	0.124	0.132	0.352	0.121	0.112
	rpl	0.064	0.066	0.049	0.061	0.033

Table 3.8s: Standard deviation of r estimates, parameter set 8, small data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.026	0.026	0.060	0.027	0.036
	b	0.068	0.067	0.264	0.068	0.075
	rpn	0.012	0.012	0.010	0.011	0.006
	al	0.061	0.063	0.184	0.057	0.067
	rpl	0.036	0.037	0.027	0.034	0.019

Key an : additive normal; b :beta; rpn : random preferences normal; al : additive normal; rpl : random preferences lognormal

Table 3.1m: Standard deviation of r estimates, parameter set 1, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.013	0.017	0.023	0.014	0.044
	<i>b</i>	0.021	0.017	0.038	0.019	0.050
	<i>rpn</i>	0.016	0.015	0.011	0.015	0.024
	<i>al</i>	0.030	0.034	0.052	0.027	0.079
	<i>rpl</i>	0.016	0.014	0.012	0.015	0.023

Table 3.2m: Standard deviation of r estimates, parameter set 2, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.007	0.009	0.011	0.007	0.026
	<i>b</i>	0.020	0.018	0.055	0.019	0.052
	<i>rpn</i>	0.008	0.007	0.006	0.007	0.012
	<i>al</i>	0.015	0.018	0.024	0.014	0.047
	<i>rpl</i>	0.009	0.008	0.006	0.009	0.013

Table 3.3m: Standard deviation of r estimates, parameter set 3, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.013	0.016	0.023	0.014	0.043
	<i>b</i>	0.019	0.017	0.037	0.019	0.048
	<i>rpn</i>	0.016	0.015	0.011	0.015	0.025
	<i>al</i>	0.030	0.035	0.054	0.029	0.083
	<i>rpl</i>	0.016	0.014	0.012	0.015	0.023

Table 3.4m: Standard deviation of r estimates, parameter set 4, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.006	0.009	0.011	0.007	0.026
	<i>b</i>	0.020	0.017	0.052	0.020	0.053
	<i>rpn</i>	0.008	0.007	0.006	0.008	0.012
	<i>al</i>	0.015	0.018	0.025	0.014	0.047
	<i>rpl</i>	0.009	0.007	0.007	0.008	0.013

Table 3.5m: Standard deviation of r estimates, parameter set 5, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.040	0.040	0.122	0.040	0.049
	<i>b</i>	0.048	0.047	0.201	0.049	0.060
	<i>rpn</i>	0.016	0.017	0.013	0.015	0.008
	<i>al</i>	0.085	0.102	0.363	0.081	0.084
	<i>rpl</i>	0.044	0.046	0.035	0.045	0.024

Table 3.6m: Standard deviation of r estimates, parameter set 6, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.020	0.019	0.058	0.020	0.032
	<i>b</i>	0.049	0.046	0.302	0.049	0.062
	<i>rpn</i>	0.008	0.009	0.007	0.008	0.005
	<i>al</i>	0.045	0.045	0.242	0.041	0.054
	<i>rpl</i>	0.026	0.026	0.020	0.026	0.014

Table 3.7m: Standard deviation of r estimates, parameter set 7, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.039	0.039	0.118	0.038	0.053
	<i>b</i>	0.049	0.047	0.176	0.049	0.061
	<i>rpn</i>	0.016	0.017	0.012	0.015	0.009
	<i>al</i>	0.085	0.097	0.312	0.082	0.086
	<i>rpl</i>	0.045	0.048	0.035	0.044	0.024

Table 3.8m: Standard deviation of r estimates, parameter set 8, medium data set

Specification		Estimated				
		<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	0.019	0.019	0.050	0.019	0.032
	<i>b</i>	0.048	0.047	0.238	0.050	0.060
	<i>rpn</i>	0.009	0.008	0.007	0.008	0.005
	<i>al</i>	0.042	0.046	0.247	0.042	0.055
	<i>rpl</i>	0.026	0.026	0.020	0.025	0.013

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 3.1I: Standard deviation of r estimates, parameter set 1, large data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.009	0.012	0.015	0.010	0.031
	b	0.014	0.012	0.025	0.014	0.037
	rpn	0.011	0.011	0.008	0.010	0.017
	al	0.021	0.024	0.040	0.020	0.057
	rpl	0.011	0.010	0.009	0.010	0.016

Table 3.3I: Standard deviation of r estimates, parameter set 3, large data set

Specification		Estimated				
		an	b	Rpn	al	rpl
True	an	0.009	0.011	0.016	0.010	0.031
	b	0.014	0.012	0.025	0.014	0.035
	rpn	0.012	0.011	0.008	0.011	0.017
	al	0.021	0.024	0.037	0.020	0.057
	rpl	0.011	0.010	0.008	0.011	0.016

Table 3.5I: Standard deviation of r estimates, parameter set 5, large data set

Specification		Estimated				
		an	b	Rpn	al	rpl
True	an	0.027	0.029	0.115	0.028	0.034
	b	0.035	0.033	0.236	0.034	0.042
	rpn	0.011	0.011	0.009	0.010	0.006
	al	0.060	0.070	0.402	0.059	0.059
	rpl	0.032	0.033	0.026	0.032	0.017

Table 3.7I: Standard deviation of r estimates, parameter set 7, large data set

Specification		Estimated				
		an	b	Rpn	al	rpl
True	an	0.027	0.027	0.145	0.028	0.037
	b	0.033	0.032	0.214	0.034	0.043
	rpn	0.011	0.011	0.009	0.011	0.006
	al	0.059	0.069	0.398	0.058	0.060
	rpl	0.031	0.034	0.025	0.030	0.016

Table 3.2I: Standard deviation of r estimates, parameter set 2, large data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.005	0.006	0.008	0.005	0.019
	b	0.014	0.013	0.076	0.014	0.036
	rpn	0.006	0.005	0.004	0.005	0.009
	al	0.011	0.013	0.017	0.010	0.033
	rpl	0.006	0.006	0.005	0.006	0.009

Table 3.4I: Standard deviation of r estimates, parameter set 4, large data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.005	0.006	0.008	0.005	0.019
	b	0.014	0.012	0.080	0.013	0.037
	rpn	0.006	0.005	0.004	0.005	0.009
	al	0.011	0.013	0.042	0.010	0.034
	rpl	0.006	0.005	0.005	0.006	0.009

Table 3.6I: Standard deviation of r estimates, parameter set 6, large data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.014	0.014	0.036	0.014	0.022
	b	0.034	0.034	0.368	0.034	0.042
	rpn	0.006	0.006	0.005	0.006	0.003
	al	0.031	0.032	0.301	0.029	0.039
	rpl	0.019	0.019	0.014	0.018	0.009

Table 3.8I: Standard deviation of r estimates, parameter set 8, large data set

Specification		Estimated				
		an	b	rpn	al	rpl
True	an	0.013	0.013	0.040	0.014	0.022
	b	0.035	0.032	0.371	0.034	0.043
	rpn	0.006	0.006	0.005	0.006	0.003
	al	0.030	0.032	0.281	0.031	0.039
	rpl	0.019	0.019	0.014	0.018	0.009

Key an : additive normal; b :beta; rpn : random preferences normal; al : additive normal; rpn : random preferences lognormal

Table 4.1s: s estimates, parameter set 1, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	27.95	56.23	5.30	49.33	1.67
	<i>b</i>	40.00	18.98	42.68	4.51	34.92	1.38
	<i>rpn</i>	10.00	31.43	91.84	10.30	60.21	2.94
	<i>al</i>	20.00	13.10	19.62	2.80	24.05	0.73
	<i>rpl</i>	2.92	31.87	96.00	10.12	60.79	3.04

Table 4.3s: s estimates, parameter set 3, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	27.89	56.39	5.26	49.19	1.65
	<i>b</i>	40.00	18.96	42.30	4.48	35.07	1.39
	<i>rpn</i>	10.00	31.52	89.56	10.36	60.21	2.97
	<i>al</i>	20.00	13.17	19.51	2.80	24.01	0.73
	<i>rpl</i>	2.92	31.90	95.70	10.02	60.49	3.01

Table 4.5s: s estimates, parameter set 5, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	26.00	136.08	1.67	44.96	0.49
	<i>b</i>	80.00	20.15	85.83	1.38	35.13	0.37
	<i>rpn</i>	10.00	91.07	1394.94	10.38	168.50	3.47
	<i>al</i>	20.00	11.59	24.63	1.10	20.59	0.19
	<i>rpl</i>	1.06	30.56	179.74	3.44	58.48	1.10

Table 4.7s: s estimates, parameter set 7, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	25.72	137.08	1.65	44.92	0.49
	<i>b</i>	80.00	20.11	87.02	1.38	35.05	0.36
	<i>rpn</i>	10.00	90.75	1390.86	10.47	168.15	3.45
	<i>al</i>	20.00	11.69	24.91	1.11	20.85	0.19
	<i>rpl</i>	1.06	30.70	177.50	3.43	58.33	1.10

Table 4.2s: s estimates, parameter set 2, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	54.06	169.37	10.48	95.47	3.21
	<i>b</i>	40.00	19.00	42.26	4.59	35.15	1.40
	<i>rpn</i>	20.00	62.33	375.08	20.69	120.33	6.18
	<i>al</i>	40.00	25.05	48.38	4.99	45.73	1.53
	<i>rpl</i>	5.41	57.18	293.34	18.68	108.81	5.60

Table 4.4s: s estimates, parameter set 4, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	54.05	165.03	10.53	95.58	3.20
	<i>b</i>	40.00	19.05	42.38	4.49	35.11	1.41
	<i>rpn</i>	20.00	62.51	384.39	20.83	119.45	6.16
	<i>al</i>	40.00	25.01	49.02	4.92	45.85	1.53
	<i>rpl</i>	5.41	56.91	291.35	18.61	107.88	5.60

Table 4.6s: s estimates, parameter set 6, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	51.55	546.74	3.17	90.36	1.00
	<i>b</i>	80.00	20.09	86.60	2.03	35.24	0.36
	<i>rpn</i>	20.00	182.72	1600.00	21.35	330.73	7.10
	<i>al</i>	40.00	22.85	109.21	2.07	41.50	0.44
	<i>rpl</i>	1.99	54.40	567.33	6.24	104.10	2.07

Table 4.8s: s estimates, parameter set 8, small data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	51.60	544.50	3.13	90.90	0.99
	<i>b</i>	80.00	20.09	85.29	2.03	35.37	0.36
	<i>rpn</i>	20.00	182.69	1600.00	21.29	329.97	7.12
	<i>al</i>	40.00	23.01	105.95	2.09	41.55	0.44
	<i>rpl</i>	1.99	54.68	557.47	6.28	104.08	2.07

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 4.1m: *s* estimates, parameter set 1, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	27.28	53.24	5.25	48.23	1.48
	<i>b</i>	40.00	18.52	41.26	3.95	34.29	1.15
	<i>rpn</i>	10.00	30.09	82.22	10.15	58.29	2.89
	<i>al</i>	20.00	12.78	18.63	2.48	23.52	0.62
	<i>rpl</i>	2.92	30.61	89.52	9.94	58.65	2.98

Table 4.2m: *s* estimates, parameter set 2, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	53.02	160.04	10.41	93.50	2.87
	<i>b</i>	40.00	18.55	41.12	4.00	34.40	1.12
	<i>rpn</i>	20.00	60.36	350.26	20.38	115.67	6.09
	<i>al</i>	40.00	24.42	46.70	4.79	44.67	1.37
	<i>rpl</i>	5.41	54.53	283.12	18.26	104.97	5.52

Table 4.3m: *s* estimates, parameter set 3, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	27.35	53.17	5.24	48.27	1.47
	<i>b</i>	40.00	18.44	41.02	3.88	34.32	1.16
	<i>rpn</i>	10.00	29.91	81.73	10.20	57.62	2.91
	<i>al</i>	20.00	12.83	18.44	2.51	23.48	0.62
	<i>rpl</i>	2.92	30.66	89.44	9.88	58.68	2.97

Table 4.4m: *s* estimates, parameter set 4, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	53.10	161.42	10.41	93.66	2.93
	<i>b</i>	40.00	18.41	41.20	4.09	34.31	1.14
	<i>rpn</i>	20.00	59.83	350.39	20.47	115.74	6.10
	<i>al</i>	40.00	24.43	46.54	4.81	44.84	1.37
	<i>rpl</i>	5.41	54.37	284.13	18.27	104.01	5.50

Table 4.5m: *s* estimates, parameter set 5, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	25.47	131.56	1.43	44.39	0.38
	<i>b</i>	80.00	19.72	82.35	1.23	34.43	0.28
	<i>rpn</i>	10.00	88.77	1402.67	10.19	167.83	3.38
	<i>al</i>	20.00	11.31	23.06	1.04	20.32	0.15
	<i>rpl</i>	1.06	29.78	166.95	3.35	57.33	1.08

Table 4.6m: *s* estimates, parameter set 6, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	50.84	530.53	2.69	89.09	0.76
	<i>b</i>	80.00	19.73	82.99	2.00	34.65	0.29
	<i>rpn</i>	20.00	178.10	1600.00	20.60	331.01	6.87
	<i>al</i>	40.00	22.44	102.50	2.02	40.67	0.35
	<i>rpl</i>	1.99	53.27	549.07	6.13	102.82	2.03

Table 4.7m: *s* estimates, parameter set 7, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	25.40	131.73	1.44	44.21	0.38
	<i>b</i>	80.00	19.61	83.56	1.23	34.35	0.28
	<i>rpn</i>	10.00	88.62	1401.38	10.19	168.71	3.38
	<i>al</i>	20.00	11.40	22.75	1.04	20.40	0.15
	<i>rpl</i>	1.06	29.82	168.26	3.36	57.55	1.08

Table 4.8m: *s* estimates, parameter set 8, medium data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	50.87	530.24	2.72	88.99	0.77
	<i>b</i>	80.00	19.82	82.69	2.00	34.57	0.29
	<i>rpn</i>	20.00	177.34	1600.00	20.65	331.73	6.90
	<i>al</i>	40.00	22.58	101.97	2.02	40.57	0.35
	<i>rpl</i>	1.99	53.37	541.65	6.14	101.78	2.02

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Table 4.1l: s estimates, parameter set 1, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	27.11	52.11	4.99	48.05	1.42
	<i>b</i>	40.00	18.29	40.82	3.46	33.99	1.08
	<i>rpn</i>	10.00	29.40	79.21	10.06	57.60	2.85
	<i>al</i>	20.00	12.64	18.11	2.27	23.37	0.59
	<i>rpl</i>	2.92	30.02	87.28	9.82	57.93	2.95

Table 4.2l: s estimates, parameter set 2, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	52.65	156.67	10.14	92.60	2.73
	<i>b</i>	40.00	18.32	40.66	3.75	34.13	1.09
	<i>rpn</i>	20.00	59.43	335.47	20.25	114.44	6.03
	<i>al</i>	40.00	24.20	45.10	4.55	44.25	1.31
	<i>rpl</i>	5.41	53.63	277.30	18.10	103.74	5.45

Table 4.3l: s estimates, parameter set 3, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	27.12	52.18	4.98	47.87	1.41
	<i>b</i>	40.00	18.31	40.54	3.48	33.84	1.09
	<i>rpn</i>	10.00	29.31	78.38	10.10	57.27	2.87
	<i>al</i>	20.00	12.62	18.03	2.30	23.37	0.60
	<i>rpl</i>	2.92	30.07	86.68	9.81	57.97	2.95

Table 4.4l: s estimates, parameter set 4, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	52.62	156.13	10.12	92.78	2.75
	<i>b</i>	40.00	18.27	40.47	3.74	33.98	1.08
	<i>rpn</i>	20.00	59.08	338.60	20.24	114.24	6.04
	<i>al</i>	40.00	24.23	45.19	4.59	44.28	1.31
	<i>rpl</i>	5.41	53.78	276.93	18.09	103.32	5.45

Table 4.5l: s estimates, parameter set 5, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	25.21	128.11	1.37	43.95	0.36
	<i>b</i>	80.00	19.56	81.49	1.16	34.22	0.26
	<i>rpn</i>	10.00	86.87	1406.74	10.12	166.22	3.35
	<i>al</i>	20.00	11.23	22.16	1.01	20.19	0.14
	<i>rpl</i>	1.06	29.42	162.04	3.32	56.66	1.07

Table 4.6l: s estimates, parameter set 6, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	50.29	522.46	2.53	88.16	0.70
	<i>b</i>	80.00	19.58	81.22	2.00	34.37	0.27
	<i>rpn</i>	20.00	174.87	1600.00	20.16	327.71	6.79
	<i>al</i>	40.00	22.28	98.52	2.00	40.29	0.33
	<i>rpl</i>	1.99	52.39	529.54	6.06	101.66	2.01

Table 4.7l: s estimates, parameter set 7, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	25.00	25.20	128.65	1.36	43.99	0.35
	<i>b</i>	80.00	19.52	81.44	1.16	34.24	0.27
	<i>rpn</i>	10.00	87.00	1407.58	10.10	167.45	3.34
	<i>al</i>	20.00	11.29	21.84	1.01	20.32	0.14
	<i>rpl</i>	1.06	29.38	161.53	3.33	56.93	1.07

Table 4.8l: s estimates, parameter set 8, large data set

		True value	Estimated value				
			<i>an</i>	<i>b</i>	<i>rpn</i>	<i>al</i>	<i>rpl</i>
True	<i>an</i>	50.00	50.43	520.29	2.53	88.30	0.69
	<i>b</i>	80.00	19.62	81.39	2.00	34.23	0.27
	<i>rpn</i>	20.00	173.88	1600.00	20.28	327.60	6.78
	<i>al</i>	40.00	22.29	98.27	2.00	40.31	0.32
	<i>rpl</i>	1.99	52.38	530.06	6.08	101.10	2.01

Key *an*: additive normal; *b*:beta; *rpn*: random preferences normal; *al*: additive normal; *rpn*: random preferences lognormal

Chapter 5

Stock Market Sectors' Dependencies Pre and Post the Financial Crisis of 2007: the cases of Greece, Italy and Portugal

5.1 Introduction

After the collapse of the fourth largest investment bank in the USA (Lehman Brothers) on 2007, which caused many markets around the world to experience a very volatile economic environment with an increased level of uncertainty (Bartram and Bodnar, 2009), the subprime crisis of the USA transmitted its shocks to the other side of the Atlantic Ocean: the European continent. The first victim of the European sovereign debt crisis was Greece in 2007-2008. Greece had a very high level of government debt. In addition to Greece, other Southern European countries, after a short time, proved that they were also in a similar financial position, and they had to deal with high debt and deficit (Samitas and Tsakalos, 2013). These European countries were named with the acronym PIIGS; this included the following countries: Greece, Italy, Ireland, Spain and Portugal. It was necessary for people to understand that this crisis was not in reality confirmed to these countries but that it was a global financial crisis, and possibly the biggest in the history of the European Union. After Greece, the next countries that experienced the deficit problem were Belgium, the UK and France.

In this chapter, we focus on the effects of the European debt crisis on the relationships of sectoral stock market indices, by testing the dependency relationships among the stock market sectors. We examine three out of the five PIIGS countries, specifically Greece, Italy and Portugal. We choose these three countries as they had the biggest drop in their average real GDP % change Year on Year after the crisis compared to Ireland and Spain (see Appendix A). We will conduct cointegration and Granger causality analysis for the stock market sector indices in each of the three countries separately. We explore these tests for the five most powerful sectors of each of the corresponding stock markets.

While there is a massive literature on the dependence relationships between countries' basic stock market indices, there are just few papers which explore the effect of the recent global financial crisis on the dependencies of the PIIGS countries' stock market main indices. More specifically, Chouliaras *et al* (2012) explore whether cointegration and causality relationships exist among the basic stock markets indices of the PIIGS countries. By using daily data from 01/02/2005 to 30/06/2011, they find long and short run relations among the main indices of the markets. Kazemi and Sohrabji (2012) look for contagion among the PIIGS: they discover higher post-crisis correlation among the PIIGS countries. Tamakoshi and Hamori (2011) look for transmission of stock indices among the European PIIGS, Germany and the UK before and during the European sovereign debt crisis. They conclude that there are short run relationships mainly from Ireland and Portugal to other countries, especially to Germany, prior to the crisis. Ahmad *et al* (2013) assess the financial contagion by looking at the contagion effect of the PIIGS countries, the USA, the UK and the Japanese markets on the BRIICKS countries' stock markets. These are Brazil, Russia, India, Indonesia, China, South Korea and South Africa. Their main findings are that Ireland, Italy and Spain have higher level of contagion than Greece for BRIICKS stock markets, while Brazil, India, Russia, China and South Africa were significantly affected by the contagion shock during the Eurozone crisis. Tiwari *et al* (2016) examine co-movements, contagion and rolling correlation between the

basic stock market indices of the PIIGS countries and Germany and the UK. Their conclusion in the short run is that correlation was high during the crisis, while in the long run the co-movements were present for the entire period.

The analysis of a market's long run and short run relationships among its sectors is very important for a country's optimal portfolio allocation and risk management. From a policy perspective, the analysis provides useful implications for the government of each country on how to take measures to rescue it when the country appears to be vulnerable. Furthermore, as Ewing (2002) states, stock sector indices act as benchmarks for the profitability of the publicly traded stocks that are included in the sector or the performance of actively managed portfolios that include publicly traded stocks of a sector. In addition, Ewing (2002) claims that indices are also used as the basis for other financial instruments, such as index mutual funds.

Based on the existing literature review on PIIGS countries presented above, there is no paper that studies the sectoral stock indices relationships of any of the PIIGS countries before and after the financial crisis. Actually, this is the first paper that tests these relationships in a country's stock market before and after a financial crisis. Furthermore, it examines these relationships in three countries which experienced the same crisis. In addition to this lack of research, we carry out this analysis by taking into consideration the benefits that a sectoral stock market analysis offers to a country's policy and its investors.

Below we refer to some of the few papers that conduct a sectoral cointegration and causality analysis. The first paper is the only one that relates to one of the countries in our analysis. Patra and Poshakwale (2008) examine the case of the Athens stock exchange using daily data from 01/01/1996 to 31/12/2003 and report that there is not a consistent cointegration relationship. However, the banking sector appears to be very influential on the return and volatility of other sectors in the short run. By using variance decomposition analysis they also show that the banking sector is able to explain 25% of the variance of the construction and

insurance sectors, as well as 15% of the variance of the industrial, investment and the holding sectors. Vardhan *et al* (2015) study the Indian sector stock price indices in the post subprime crisis period and report minimal benefits from diversifying investments to different sectors. Here again, the banking sector plays a key role and seems to move other indices, while other indices seem to be driven by the Realty and the Metal sectors. Ewing (2002) studies the interrelationships of five important S&P stock indices and how shocks of each of the indices are transmitted to other indices. Ahmed (2012) analyses the interdependence of the four market sector indices of the Qatar Exchange. His results suggest that there is a cointegration relationship and that the banking and financial institutions sector seems to drive other sectors in the short run. Wang *et al* (2005) study the dependence relationships across and within the two Chinese Stock Exchanges in Shanghai and Shenzhen and show a high degree of interdependence among the sectors. This means that a portfolio based only on the sectors of the Chinese stock markets offers limited diversification. Policymakers could take this into consideration, and possibly try to create policies to stop any negative transmission of shocks between the sectors. Arbeláez *et al* (2001) assesses the linkages among the sectoral stock indices of the former Medellin Stock Exchange in Colombia. The results show that indices exhibit long-term relationships between them and short-term linkages in half of the cases. In addition, impulse response results show that responses to innovations in other indexes are small but also fast and persistent.

5.2 Data

In our analysis, we investigate the interdependencies among the five major stock market sector indices in each of the PIG (Portugal, Italy and Greece) countries: Athens Stock Market Sector Indices, Italy Stock Market Sector Indices and Portugal Stock Market Sector Indices. The sample consists of the natural logarithm of monthly closing prices for each sector index

from the 1st of June 1998 to the 1st of November 2016. This period is divided into two sub-periods: 01/06/1998 to 01/08/2007 and 01/09/2007 to 01/11/2016. The data was obtained from the Datastream database. We consider the first sub-period as the pre-crisis period, and the second sub-period as the post-crisis period; there are 111 observations in each. We use the sector classification ICB of Datastream which includes the following ten sectors: Basic Materials, Consumer Goods, Consumer Services, Financials, Health Care, Industrials, Oil and Gas, Technology, Telecommunications and Utilities. After removing the sectors for which we do not have a full data sample for the period under investigation, we select the five main sectors of each country based on their share of the country's total stock market capitalization. Based on this analysis, we use the Financials, Telecommunications, Consumer Services, Consumer Goods, and Basic Materials sectors for Greece, the Financials, Oil and Gas, Utilities, Telecommunications, and Consumer Goods sectors for Italy, and the Financials, Utilities, Consumer Services, Telecommunications and Industrials sectors for Portugal. For the specific percentage of each sector within the country's total stock market capitalization, please see Appendix B.

5.3 Methodology

The first step in the data analysis is to check the time series for stationarity. A time series is called stationary when its main characteristics are time-invariant; that is, when its mean, variance and covariance are constant over time. If this is not the case, the time series is said to be non-stationary. When a time series V_t is not stationary in levels but its p^{th} differences are stationary, it is called integrated of order p – abbreviated as $I(p)$. For instance, if V_t is non-stationary in level prices but it is stationary in 1st differences, V_t is integrated of order 1 with abbreviation $I(1)$, while the abbreviation of its first differences is $I(0)$. Let us suppose that the time series V_t is not stationary. In this case there are some problems. The first is that

the results we obtain from the classical regression techniques – such as the OLS estimators – are not reliable. That is, the dependency results of the tests can be spurious, meaning that it might not exist in reality. The second problem is that the V_t suffers from permanent effects. That is, the effects of an unpredictable sudden change (shock) on a stationary time series will be gradually reduced to zero. On the other hand, the effects of a shock on a non-stationary time series last forever. Moreover, a stationary time series is also called mean-reverting because its values are always fluctuating around the underlying mean, but the values of the non-stationary V_t can increase or decrease continuously following an unpredictable path, such as a random walk. A random walk without a drift (intercept) is the most well-known example of a process that a non-stationary time series can follow. This model suggests that the value of V is equal to its value yesterday plus an unpredictable shock (ε_t). It is an autoregressive model of order one (AR(1)) with the autoregressive parameter equal to unity and it is given as follow:

$$\text{Random walk without a drift:} \quad V_t = aV_{t-1} + \varepsilon_t \quad (1)$$

where $a = 1$. In order for V_t to be stationary, a (the coefficient of V_{t-1}) in absolute value should be less than one. In the case that V_t is non-stationary, it is also said that V_t has a unit root as a will be equal to one. This can be explained with the help of the lag operator (L) in equation (1):

$$V_t = aV_{t-1} + \varepsilon_t \Leftrightarrow L^0V = aL^1V + \varepsilon_t \Leftrightarrow V - aLV = \varepsilon_t \Leftrightarrow V(1 - aL) = \varepsilon_t \quad (2)$$

$(1 - aL)$ is called the characteristic equation and the value of L that sets $(1 - aL)$ to zero is called the characteristic root; that is $(1 - aL) = 0$ and $L = \frac{1}{a}$. The value of a is on the unit circle, if its value is equal to 1 the time series has a unit root so it is non-stationary, while when a 's absolute value is less than 1, the time series does not have a unit root and it is stationary.

5.3.1 Dickey-Fuller GLS unit root test

The most common unit root test is the Augmented Dickey-Fuller (ADF) unit root test proposed by Dickey and Fuller (1979). This test is based on a modified version of equation (1), which is produced by subtracting the value of V at the $t - 1$ period, adding q lagged values of the first differences of V and adding the X_t vector which contains the possible exogenous data, and its coefficient δ :

$$\begin{aligned}\Delta V_t &= aV_{t-1} - V_{t-1} + \delta X_t + \sum_{i=1}^q \gamma_i \Delta V_{t-i} + \varepsilon_t = (a - 1)V_{t-1} + \delta X_t + \sum_{i=1}^q \gamma_i \Delta V_{t-i} + \varepsilon_t \\ &= \beta V_{t-1} + \delta X_t + \sum_{i=1}^q \gamma_i \Delta V_{t-i} + \varepsilon_t\end{aligned}\quad (3)$$

In equation (3) β is equal to $(a - 1)$ and ε_t is the disturbance term (errors) which is assumed to be white noise: that is, its mean is zero and its variance is constant over time, and is abbreviated as $\varepsilon_t \sim IID(0, \sigma_\varepsilon^2)$. *IID* means that these random variables (shocks) ε_t are Independent and Identically Distributed. Moreover, the X_t vector can contain either an intercept and a trend, or only an intercept, or neither. Thus the ADF unit root test relies on three different models which depend on the exogenous data of the X_t vector. The ADF unit root test has the following null and alternative hypotheses:

(Null Hypothesis) $H_0: \beta = 0 \Leftrightarrow Y_t$ has a unit root

(Alternative Hypothesis) $H_1: \beta < 0 \Leftrightarrow Y_t$ does not have a unit root

In order to decide if the null hypothesis should be rejected or not, the t-statistic of β should be computed based on the following formula below where $\hat{\beta}$ is the estimator of β and $se(\hat{\beta})$ is the standard error of the $\hat{\beta}$ estimator.

$$\hat{t} = \frac{\hat{\beta}}{se(\hat{\beta})}\quad (4)$$

The value of $\hat{\tau}$ is then compared to the Dickey–Fuller critical values and the null hypothesis is rejected if $\hat{\tau}$ in absolute value is larger than the absolute value of the critical value at the selected significance level, here at 5%³¹. Note that there are different critical values for each of the three ADF models.

Although the ADF unit root test is a very common tool to check stationarity, it is not the preferred one when it comes to small datasets as it is known to have low power against the alternative hypothesis; that is, it tends to not reject the null hypothesis of unit root existence more often than it should. As the number of observations in this study is 111 per period tested, a modified version of the ADF unit root test, called the Dickey–Fuller GLS unit root test (DF-GLS) is used (Elliot *et al*, 1996). The DF-GLS unit root test is based on the de-trended V_t data, denoted by V_t^d which are equal to the V_t data minus the exogenous variables X_t . Specifically, the method suggested by Elliot *et al* starts by defining the value of V_t based on the time t and the value of β .

$$d(V_t|\beta) = \begin{cases} V_t & , \text{ if } t = 1 \\ V_t - \beta V_{t-1} & , \text{ if } t > 1 \end{cases}$$

Then δ is estimated by the following model where X_t is either an intercept (1) or an intercept and a trend (1+t):

$$d(V_t|\beta) = d(X_t|\beta)\delta(\beta) + u_t \tag{5}$$

$$X_t = \begin{cases} 1 \\ 1 + t \end{cases}$$

In order to run the model in equation (5), β should be defined. Elliot *et al* (1996) suggested using the following $\bar{\beta}$ which depends upon the value of X_t :

$$\bar{\beta} = \begin{cases} 1 - 7T & \text{if } X_t = \{1\} \\ 1 - 13.5T & \text{if } X_t = \{1 + t\} \end{cases}$$

³¹ We choose the 5% significance level, as this is the significance level that most researchers use. The 1% significance level is considered very strict and the 10% significance level too loose.

Finally, the V_t and ΔV_t in equation (3) are replaced by the de-trended values $V_t^d = V_t - X_t\delta(\beta)$ and $\Delta V_t^d = V_t - V_t^d$ respectively. Thus, the DF-GLS model is as follow:

$$\Delta V_t^d = \beta V_{t-1}^d + \sum_{i=1}^q \gamma_i \Delta V_{t-i}^d + \varepsilon_t \quad (6)$$

In the software we use, Eviews, the number of lags (q) is defined by using an information criterion such as the Schwarz Information Criterion and the Akaike Information Criterion. The null and the alternative hypotheses of the DF-GLS test remain the same with the ones in ADF unit root test, as well as the formula for the t-statistic.

(Null Hypothesis) $H_0: \beta = 0 \Leftrightarrow Y_t$ has a unit root

(Alternative Hypothesis) $H_1: \beta < 0 \Leftrightarrow Y_t$ does not have a unit root

$$\hat{t} = \frac{\hat{\beta}}{se(\hat{\beta})}$$

However, the value of this statistic is compared with different critical values (cv_{GLS}) than those used in the ADF unit root test. In the case that $X_t = \{1\}$ (only intercept), the critical values of the DF-GLS test are same as the ones in the ADF model without intercept and trend. However, in the case that $X_t = \{1 + t\}$ (intercept and trend) Elliot *et al* (1996) have estimated different critical values (Appendix C – Table 1). If $|\hat{t}| > |cv_{GLS}|$, the null hypothesis of unit root (non-stationarity) is rejected. If the inequality is reversed, the null hypothesis cannot be rejected and the time series V_t is declared non-stationary.

5.3.2 The Johansen cointegration test

Assuming that the variables of interest are I(1) the next step is to analyse the data for long-run relationships by applying a cointegration test. There are many different cointegration techniques widely used when researchers want to examine the possibility of two or more variables to be cointegrated (have a long-run equilibrium relationship). Namely, two time

series are called cointegrated when their price deviations are generally nearly the same by size and direction (positive or negative), so that their value paths have an almost stable distance. A common requirement for all the different cointegration tests is that the dependent variables are integrated of the same order (usually integrated of order one). After testing all the time series for unit root, the Johansen cointegration test is in order (Johansen, 1991). The Johansen cointegration test is based on a VAR model of order q , where V_t is a $n \times 1$ vector which contains the non-stationary time series which we want to test for cointegration, q is the number of lags, X_t is a vector consisting of the possible deterministic variables (intercept and trend) and ε_t is *IID*.

$$V_t = K_1 V_{t-1} + K_2 V_{t-2} + \dots + K_q V_{t-q} + \Gamma X_t + \varepsilon_t \quad (7)$$

This model can be rewritten in error correction form with dependent variable a vector consisting of the stationary 1st differences of the time series tested as follows:

$$\Delta V_t = \Pi V_{t-1} + \sum_{i=1}^{q-1} A_i \Delta V_{t-i} + \Gamma X_t + \varepsilon_t \quad (8)$$

with the coefficient matrix $\Pi = \sum_{i=1}^q K_i - I$ and $A_i = -\sum_{j=i+1}^q K_j$. I is the identity matrix and the number of lags is defined through an unrestricted VAR model on V_t including an intercept and a trend if your variables are trended and based on the Akaike Information Criterion (AIC), so that there is no serial correlation among the residuals. A proof of how equation (8) can give the same result as equation (7) can be found on Appendix D. The formula for the AIC is as follows:

$$AIC = 2\theta - 2\ln(\hat{L})$$

where θ is the number of the parameters to be estimated and \hat{L} is the maximised value of the likelihood function of the model in equation (8).

If the rank (r) of the coefficient matrix Π is less than the number of the endogenous variables n ($r < n$) and for suitable $n \times r$ matrices α and β , $\Pi = \alpha\beta'$ and $\beta'V_t \sim I(0)$. r is the number of

the cointegration relations (rank), β is a matrix with cointegration vectors in each column, β' is the transpose matrix of β and the data of matrix α are the adjustment parameters which will be used in a Vector Error Correction model in the case that cointegration exists. When r is equal to zero there is no cointegration among the non-stationary time series, while when r is between 1 and $n-1$ there are r cointegration equations among them. In the case that the rank of Π matrix is full ($r=n$), the logarithmic prices V_t are stationary.

The first thing in the Johansen cointegration test is to estimate the matrix Π and test a series of null hypotheses for $r = 0, 1, 2, 3, \dots, n-1$ against an alternative. Johansen (1991) recommended two different cointegration tests to choose from. The first is the Trace Test and the second is the Maximum Eigenvalue Test with the following formulas, where λ_j is the j^{th} largest eigenvalue of the coefficient matrix Π .

$$\lambda_{trace} = -T \sum_{j=r+1}^n \log(1 - \hat{\lambda}_j) \quad (9)$$

$$\lambda_{max} = -T \sum_{j=r+1}^n \log(1 - \widehat{\lambda}_{r+1}) \quad (10)$$

The alternative hypotheses are $r=n$ and $r=r+1$, respectively. The process of the tests is to test sequentially the null hypotheses of $r = 0, 1, \dots, n-1$, until we fail to reject the null. The null hypothesis is rejected when the corresponding test statistic is larger than the matching critical value at the 5% significance level. Note that each null hypothesis has different critical values which depend on the existence of deterministic variables and the number of dependent variables in the equation (8).

5.3.3 Granger causality test

Apart from the long-run relationships between the data, the short-term interactions between them can be tested by using the methodology that Granger (1969) suggested. The following bivariate VAR model is applied

$$\begin{pmatrix} \Delta V_1 \\ \Delta V_2 \end{pmatrix} = \begin{pmatrix} \gamma_0 \\ \varphi_0 \end{pmatrix} + \begin{pmatrix} \delta_1 \\ \delta_2 \end{pmatrix} (V_{2,t-i} - V_{1,t-i}) + \begin{pmatrix} \sum_{i=1}^q \gamma_{1,i} \\ \sum_{i=1}^q \varphi_{1,i} \end{pmatrix} \Delta V_{1,t-i} + \begin{pmatrix} \sum_{i=1}^q \gamma_{2,i} \\ \sum_{i=1}^q \varphi_{2,i} \end{pmatrix} \Delta V_{2,t-i} + \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix}$$

$$\Leftrightarrow \Delta V_1 = \gamma_0 + \delta_1(V_{2,t-i} - V_{1,t-i}) + \sum_{i=1}^q \gamma_{1,i} \Delta V_{1,t-i} + \sum_{i=1}^q \gamma_{2,i} \Delta V_{2,t-i} + \varepsilon_{1,t} \quad (11)$$

$$\text{and } \Delta V_2 = \varphi_0 + \delta_2(V_{2,t-i} - V_{1,t-i}) + \sum_{i=1}^q \varphi_{1,i} \Delta V_{1,t-i} + \sum_{i=1}^q \varphi_{2,i} \Delta V_{2,t-i} + \varepsilon_{2,t} \quad (12)$$

There are two different versions of equations (11) and (12) depending on whether a cointegration relationship exists or not among the data. Equations (11) and (12) are used when the time series tested are found to be cointegrated, while in the case of no cointegration $\delta_1 = \delta_2 = 0$ so $\delta_1(V_{2,t-i} - V_{1,t-i})$ and $\delta_2(V_{2,t-i} - V_{1,t-i})$ are not included. This is because when there is cointegration, a Vector Error Correction Model (VECM) should be used with δ_1 and δ_2 representing the speed of adjustment to the equilibrium in the case of cointegration. The null and alternative hypotheses of equation (11) are as follow

Cointegration found

H₀: $\gamma_{2,1} = \gamma_{2,2} = \dots = \gamma_{2,q} = 0$ and $\delta_1=0$ (V_2 does not Granger cause V_1)

H₁: at least one of the $\gamma_{2,i}$ and δ_1 to be different from zero (V_2 Granger causes V_1)

No cointegration found

H₀: $\gamma_{2,1} = \gamma_{2,2} = \dots = \gamma_{2,q} = 0$ (V_2 does not Granger cause V_1)

H₁: at least one of the $\gamma_{2,i}$ to be different from zero (V_2 Granger causes V_1)

The same hypotheses, but for the $\varphi_{1,i}$ coefficients of $\Delta V_{1,t-i}$ and δ_2 , are tested for the equation (12). The no rejection of the null hypothesis imposes that the values of one variable in the test do not influence the future values of the other. Specifically, if V_2 variable is found to Granger cause the V_1 variable, it means that the past values of V_2 variable contain useful information for the prediction of the future values of V_1 variable.

5.3.4 Generalised impulse responses and variance decomposition

Researchers usually use Cholesky factorisation to orthogonalise the VAR innovations (shocks), so that they are uncorrelated contemporaneously. However, this method can be complicated because of its high level of sensitivity to the variable ordering. That is, a small change in the order of the variables in the VAR model can produce extremely different results. For this reason we use the generalised impulse response analysis (Koop *et al*, 1996, Pesaran and Shin, 1998). Generalised Impulse Response Functions (IRF) analysis is invariant of the ordering of the data series in the VAR or the VECM model.

Let us refer to equation (8). ΔV_t can be expressed as a moving average representation as below:

$$\Delta V_t = \sum_{i=1}^{\infty} a_i \varepsilon_{t-i} \quad , \text{ for } i=1, \dots, \infty \text{ and } t=1, 2, 3, \dots, T$$

Where a_i are the impulse response functions. The generalised formula of them is given by the equation:

$$\Psi_j(n) = \sigma_{jj}^{-\frac{1}{2}} \alpha_n \Sigma e_j \quad , \text{ for } j=0, 1, 2, \dots, n$$

Where Σ is the variance-covariance matrix with dimensions $n \times n$, jj is the jj^{th} element of the Σ matrix, and e_j is a $n \times 1$ vector with unity as its j^{th} row and zero elsewhere. The above equation is also equal to the difference between the expected value of V_t when we know the information set of yesterday (Ω_{t-1}) and the size of the shock (μ), and the expected value of the V_t when we know only the Ω_{t-1} :

$$\Psi_j(n, \mu, \Omega_{t-1}) = E(V_{t+1} | \varepsilon_t = \mu, \Omega_{t-1}) - E(V_{t+1} | \Omega_{t-1})$$

For the variance decomposition we work with the following process (Sheng and Tu, 2000).

We start the analysis from the equation (7)

$$V_t = K_1 V_{t-1} + K_2 V_{t-2} + \dots + K_q V_{t-q} + \Gamma X_t + \varepsilon_t$$

which can be rewritten as a vector moving average (VMA) process

$$V_t = \mu + \sum_{q=0}^{\infty} K_1^q \varepsilon_{t-q}$$

where $\mu = (\mu + K_1 + K_2 + \dots + K_q)$ and K_0 is the unconditional mean of V_t . By using the above equation to conditionally forecast the n-step ahead V_{t+n} we get

$$V_{t+n} - E_t V_{t+n} = \sum_{q=0}^{n-1} K_1^q \varepsilon_{t-q}$$

where E_t is the conditional expectations operator. The forecast error variance of the first sector in our 5-sector system ($V_{1,t}$) is given below by

$$\sigma_1(n)^2 = \sigma_1^2 \sum_{q=1}^{n-1} k_{1,2}(q)^2 + \dots + \sigma_5^2 \sum_{q=1}^{n-1} k_{1,5}(q)^2$$

where K_1^q is a matrix 5x5 which contains the values of k_{ij} for two sectors i and j of the system and $\text{Var}(\varepsilon_{ij}) = \sigma_i^2$. Thus the ratio which gives the percentage by which the forecast error variance of the first sector in the system is due to innovations in another sector of the system (p) is given by

$$W_1(p) = \sigma_p^2 \sum_{j=1}^{n-1} \alpha_{1,p}(j)^2 / \sigma_1(n)^2$$

This process was frequently used in the past. However, during recent years the generalised formula of forecast error variance is more frequently used. The reason is that the method we use depends on the ordering of the variables of the system. Namely, the results of the test is possible to change significantly if we change the ordering of the variables in the VAR model. On the other hand, the generalised version of the variance decomposition is independent of the variables ordering. In this study, we order the sectors in the VAR models

based on their average correlation with the rest of the sectors from the one with the lowest average correlation towards the one with the greatest average correlation.

5.3.5 Dynamic conditional correlation (DCC)

Correlation tests examine whether the information that one variable contains can help to predict how another variable will act. The correlation coefficient of two variables V_i and V_j is abbreviated as ρ_{ij} and its range is between -1 and 1. In the case that $\rho_{ij} = -1$, there is a perfect negative correlation between V_i and V_j which means that if the price of V_i drops by 2% then the price of V_j will increase by 2%. In the opposite case that $\rho_{ij} = 1$, it is said that V_i and V_j are perfectly positive correlated which means that if the price of one of them decreases by 2%, then the price of the other one will also decrease by 2%. Finally in the case where $\rho_{ij} = 0$, the two variables are not correlated at all. The correlation coefficient can also be equal to any decimal number between -1 and 1, and it is said that when ρ_{ij} is closer to 0 there is weak correlation, while when ρ_{ij} is closer to -1 and 1 there is high-degree of negative and positive correlation, respectively. Correlation is a very important tool especially for investors, as it can show if there are diversification benefits among the assets. For instance, if V_i and V_j represent two stocks and they have ρ_{ij} equal to -1 or a decimal number close to -1, portfolio investors can see that there are diversification benefits as a big loss on V_i will be balanced by a gain in V_j . This is the reason that many investors build portfolios with negatively correlated international stocks, as these stocks are exposed to different countries, market economies and laws.

As the economic environment continuously changes, it is expected that the correlation between any two assets will also potentially change over time. Engle (2002) proposed the dynamic conditional correlation model (DCC) which has the flexibility of univariate GARCH

models coupled with parsimonious parametric models for the correlations. Consider the $n \times 1$ vector of logarithmic returns ΔV_t of n assets at time t , which is given by

$$\Delta V_t = E(\Delta V_t / I_{t-1}) + \pi_t \quad (13)$$

where $E(\Delta V_t / I_{t-1})$ is the $n \times 1$ vector of the expected value of the conditional ΔV_t with I_{t-1} to represent all the past information, and $\pi_t = H^{1/2} v_t$ with zero mean and conditional variance H_t , which is the $n \times n$ matrix of conditional variances of π_t at time t . Also v_t is a $n \times 1$ vector of IID errors such that its mean is zero and its variance is the identity matrix of order n , (I_n). Moreover, $H_t = D_t R_t D_t$ where D_t is the $n \times n$ diagonal matrix of the conditional standard deviations of π_t at time t given by

$$D_t = \text{diag} \left(h_{1,t}^{1/2} \dots h_{n,t}^{1/2} \right)$$

where $h_{i,t}$ follows a GARCH(1,1) model as follow

$$h_{i,t} = \gamma_{i,0} + \alpha_{i,1} \pi_{i,t-1}^2 + \beta_{i,1} h_{i,t-1}.$$

R_t is the $n \times n$ conditional correlation matrix with elements equal to or less than one, given by

$$R_t = \text{diag} \left(q_{11,t}^{-1/2} \dots q_{nn,t}^{-1/2} \right) Q_t \text{diag} \left(q_{11,t}^{-1/2} \dots q_{nn,t}^{-1/2} \right) \quad (14)$$

where $Q_t = (q_{ij,t})$ and it is a $n \times n$ symmetric positive definite matrix with elements given by

$$q_{ij,t} = (1 - \alpha - \beta) \bar{p}_{ij} + \alpha \pi_{i,t-1} \pi_{j,t-1} + \beta q_{ij,t-1} \quad (15)$$

where \bar{p}_{ij} is the unconditional variances of π_t . H_t should be positive definite as it is a covariance matrix, so R_t and D_t should be positive definite as well. Also α and β should be positive and $\alpha + \beta$ should be less than 1 in order to ensure that the H_t will be positive definite. The above model (DCC) is a generalisation of the Constant Conditional Correlation (CCC) model by Bollerslev (1990). The difference between the DCC model and the CCC model is that the latter considers the correlation matrix (R) to be time-invariant, while the former considers it to change over time.

5.3.6 Time-varying volatility

Finally, we will check if the volatilities of the time series are unstable over time: that is, if they are time-varying. To do so, we follow the procedure suggested by Griffiths *et al* (2008) and we apply an Autoregressive Conditional Heteroskedastic (ARCH) model of order p [ARCH(p)]. By regressing the data (ΔV_t) on an intercept (c) and the random error terms ε_t (also called innovations),

$$\Delta V_t = c + \varepsilon_t \quad (16)$$

we obtain the residuals ($\hat{\varepsilon}_t$) which are used on the ARCH model. ε_t are normally distributed with zero mean and h_t variance which follows an ARCH process. Let us suppose that the order of this model is one and it is as follows:

$$h_t = \alpha_0 + a_1 \varepsilon_{t-1}^2 \quad (17)$$

There are two restrictions about the values of α_0 and a_1 . Namely, α_0 and a_1 should be positive in order to ensure that the variance will be positive, and a_1 should be less than 1 because otherwise the variance will continuously increase.

In order to ensure that an ARCH(1) model is a proper one for the data, the estimated $\hat{\varepsilon}_t^2$ should be regressed on their lag values $\hat{\varepsilon}_{t-1}^2$ and then check the existence of ARCH effects through an ARCH LM test.

$$\hat{\varepsilon}_t^2 = \beta_0 + \beta_1 \hat{\varepsilon}_{t-1}^2 + \eta_t \quad (18)$$

In order to examine if there are ARCH effects in the residuals, the null of $H_0: \beta_1 = 0$ is tested against the alternative $\beta_1 \neq 0$. If the null hypothesis is not rejected, there are no ARCH effects and the ARCH model of order 1 is a good fit for the data. In the opposite case, there are ARCH effects and the order of the model, p , should be increased by 1. The value of p is increased until the null hypothesis of no ARCH effect is not rejected. When the order of the

ARCH model is specified, we generate the ARCH-Variance Series and their plot gives the time-varying volatility charts (Engle *et al*, 1991).

To sum up, we use a variety of tests to identify the interdependency relationships among the sectors of Greece, Italy and Portugal. We begin with the descriptive statistics to see how the statistical characteristics of the sectoral returns changed after the financial crisis of 2008. Then, in order to have consistent and reliable results we use the Dickey–Fuller GLS test to examine the sectors for stationarity and we apply the Johansen cointegration test to the sectors that are integrated of the same order to check the possibility of a long-run equilibrium among the sectors. Thereafter, we use an unrestricted VAR model for the non-cointegrated sectors and a VECM for the cointegrated ones to get the results of Granger causality test (short-run dynamics), the impulse responses and the variance decomposition. From the impulse responses we check if the sectors react positively or negatively to a positive shock in another sector of the system and how long it takes for the effects of this shock to die out, while from the variance decomposition we obtain the percentage of a change in the variance of a sector which is due to its own innovations. Finally, we produce the time-varying volatilities of the sectors and the time-varying correlations between them to observe how and by how much they are affected by the financial crisis of 2008.

5.4 Empirical results

5.4.1 Greece

5.4.1.1 Descriptive statistics

Table 1 displays a summary of the main statistical characteristics for the logarithmic returns of the five most important market economic sectors of Greece: basic materials (BASMAT), consumer goods (CONGDS), consumer services (CONSVS), financials (FIN) and

telecommunications (TEL). Panel A contains the results for the whole period under investigation while Panels B and C exhibit the results for the pre-crisis period and post-crisis period respectively. There are several interesting results that come out of this table. The first thing to be noticed is that during the pre-crisis period all the average returns of the sectors are positive, while during the post-crisis period they are negative except that of CONGDS which is reduced by almost 100% but it remains positive. The average returns of the five sectors for the whole period are also negative except CONGDS which is positive. This similarity to the post-crisis period results shows that most of the returns were strongly negative during the post-crisis period. This is a result that supports the fact that Greece is still struggling with the financial crisis. Another important outcome is that the FIN sector is the one that experienced the largest decrease in its average return of approximately 1064%. This result illustrates the high degree of financial crisis that Greece faced during the last nine years. The average returns of the remaining sectors also decreased by more than 100%.

By comparing the standard deviation results, which indicate the risk that these sectors have during the different periods, it can be detected that during the pre-crisis period the CONGDS has the highest risk but also the highest return which is in line with the Markovitz theory that a higher return comes with higher risk. However, during the post-crisis period this theory does not hold as the CONGDS sector, which has the highest return, also has the lowest risk; and the FIN sector, which has the lowest return, has the largest risk; with big differences in the other sectors, something that also holds for the whole period. Moreover, the majority of the sectors do not follow a normal distribution as the Jarque-Bera test of normality is rejected, which means that the null hypothesis of skewness equal to zero and kurtosis equal to three is not accepted.

Table 1

Descriptive statistics of returns in different periods – Greece

Period	BASMAT	CONGDS	CONSVS	FIN	TEL
<i>Panel A Whole Period 01/06/1998 – 01/11/2016</i>					
Mean	-0.0012	0.0090	-0.0006	-0.0287	-0.0030
Median	-0.0060	0.0161	0.0041	-0.0151	-0.0025
Maximum	0.3249	0.4685	0.5162	0.5155	0.5193
Minimum	-0.2652	-0.6879	-0.3605	-0.9494	-0.7302
Std. deviation	0.0993	0.1308	0.0994	0.1645	0.1166
Skewness	0.3618	-0.5975	0.4649	-0.9954	-0.5479
Kurtosis	3.9263	7.1626	6.9709	8.1442	11.0205
Jarque-Bera	12.723***	172.71***	153.16***	280.17***	603.42***
<i>Panel B Pre-crisis Period 01/06/1998 – 01/08/2007</i>					
Mean	0.0081	0.0166	0.0073	0.0067	0.0015
Median	-0.0019	0.0284	0.0071	0.0009	0.0011
Maximum	0.3249	0.4685	0.5162	0.2223	0.2744
Minimum	-0.2051	-0.6879	-0.2739	-0.3275	-0.2040
Std. deviation	0.0893	0.1673	0.1112	0.0930	0.0776
Skewness	0.7063	-0.5593	0.8960	-0.2510	0.1710
Kurtosis	4.6761	5.1666	6.9610	4.1980	4.4391
Jarque-Bera	22.0212***	27.2506***	86.6272***	7.7326**	10.0279***
<i>Panel C Post-crisis Period 01/09/2007 – 01/11/2016</i>					
Mean	-0.0099	0.0008	-0.0085	-0.0642	-0.0081
Median	-0.0123	0.0097	0.0041	-0.0638	-0.0092
Maximum	0.2735	0.1491	0.2073	0.5155	0.5193
Minimum	-0.2652	-0.3584	-0.3605	-0.9494	-0.7302
Std. deviation	0.1082	0.0799	0.0864	0.2084	0.1460
Skewness	0.2370	-1.1214	-0.6564	-0.5564	-0.5113
Kurtosis	3.3348	6.1696	4.6233	5.6638	8.6379
Jarque-Bera	1.5432	69.1020***	19.9771***	38.1978***	150.480***

Notes: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. The summary statistics represent the five most important market economic sectors of Greece: Basic Materials (BASMAT), Consumer Goods (CONGDS), Consumer Services (CONSVS), Financials (FIN) and Telecommunications (TEL). The null hypothesis of the Jarque-Bera test is the skewness to be equal to zero and the kurtosis to be equal to three.

5.4.1.2 Dickey-Fuller GLS unit root test

The Dickey-Fuller GLS test is used in order to check the existence of a unit root in the levels data as they should be integrated of the same order for the possibility of the existence of a long-run equilibrium relationship among them (cointegration). After visual inspection of the

graphs, we found that nearly all the Greek market sectors data do not exhibit a trend. However, the FIN sector during the post-crisis period does exhibit a trend, hence a trend is used in the Dickey-Fuller GLS test. Table 2 displays the results of the Dickey-Fuller GLS for both the pre-crisis period (Panel A) and the post-crisis period (Panel B), where the null hypothesis of unit root in levels is not rejected for any of the sectors in the two sub-periods as the t-statistic is much smaller than the critical value at the 5% significance level. However, the null hypothesis of unit root in returns (first differences) is strongly rejected at 1% for all the sectors and sub-periods as the t-statistic is much larger than the critical value at 1%. Thus, the DF-GLS unit root test gives the result that the price level of all the five sectors is integrated of order one (I(1)) which means that the price level has a unit root (that is, it is non-stationary) while the returns are found to be I(0), that is stationary. These result suggests that these five sectors can be tested for cointegration before and after the financial crisis.

Table 2
Dickey-Fuller GLS unit root test - Greece

Period		BASMAT	CONGDS	CONSVS	FIN	TEL
<i>Panel A – Pre-crisis period</i>						
<i>Log prices</i>						
	Lag	0	0	0	0	0
	τ_{μ}	-0.3127	-0.2761	-0.8945	-0.6865	-1.2710
<i>Log returns</i>						
	Lag	2	2	0	0	0
	τ_{μ}	-2.9147***	-2.9603***	-8.9144***	-9.6197***	-11.014***
<i>Panel B – Post-crisis period</i>						
<i>Log prices</i>						
	Lag	0	0	0	0	0
	τ_{μ}	-0.5221	-1.2606	-0.4100	-1.6282	-0.8408
<i>Log returns</i>						
	Lag	0	2	0	0	2
	τ_{μ}	-7.2189***	-2.8492***	-9.6554***	-9.3712 ***	-3.6445***

Notes: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. The shaded cells indicate trend existence. The critical values for Dickey-Fuller GLS unit root test when there is no trend are -2.586350, -1.943796 and -1.614784 at 1%, 5% and 10% levels of significance, respectively. In the case of trend existence, the critical values change to -3.568000, -3.020000 and -2.730000, respectively.

5.4.1.3 Johansen cointegration test

Table 3 illustrates the results of the multivariate Johansen cointegration test. This examines the possibility of the existence of a long-run equilibrium relationship among these five sectors. In order to apply this test the number of lags that will be used should be defined. We use an unrestricted VAR model on the level of the series together with the Akaike Information Criterion (AIC) to determine the number of lags needed; AIC suggests one lag. In order to ensure that this number of lags is enough, we check the residuals of the unrestricted VAR model for serial correlation to confirm that one lag is enough to eliminate the existence of autocorrelation in the data. The null hypothesis of the serial correlation test is the non-existence of autocorrelation in the residuals of the VAR model against the alternative of existence. The null hypothesis of the Johansen cointegration test changes every time as it can be seen from the first column of Table 3.

The first null hypothesis tests the non-existence of a cointegration relationship which is actually not rejected in this case for both pre-crisis and post-crisis periods. This means that there is no long-run relationship among the price series of the Greek sectors for both sub-periods and so there cannot exist long-run causality relationships as well. In order to reject the null hypothesis the trace estimator or the eigenvalue estimator should be larger than the matching critical values at the 5% significance level.

Table 3
Multivariate Johansen cointegration test - Greece

	Pre-crisis		Post-crisis	
	λ_{trace}	λ_{max}	λ_{trace}	λ_{max}
$r = 0$	67.7566	30.6184	61.0101	29.4158
$r \leq 1$	37.1383	16.3510	31.5943	15.2595
$r \leq 2$	20.7873	11.8894	16.3348	10.0195
$r \leq 3$	8.8978	6.8297	6.3153	5.4361
$r \leq 4$	2.0682	2.0682	0.8793	0.8793

Notes: The optimal lag length for the cointegration testing is defined by the Schwarz information criterion (SC). ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table 4
Granger causality test results – Greece

Panel A: Pairwise Granger causality results

	<u>Pre-crisis period</u>	<u>Post-crisis period</u>
BASMAT→CONGDS	0.583556	3.404276*
CONGDS→BASMAT	0.352693	1.270877
BASMAT→CONSVS	0.276462	1.638332
CONSVS→BASMAT	0.247723	0.525996
BASMAT→FIN	0.214531	0.372657
FIN→BASMAT	0.076264	0.030619
BASMAT→TEL	0.642589	0.994387
TEL→BASMAT	1.044656	0.010374
CONGDS→CONSVS	0.021123	0.238803
CONSVS→CONGDS	0.158511	0.183534
CONGDS→FIN	0.652362	0.136352
FIN→CONGDS	2.914622*	0.794077
CONGDS→TEL	0.056950	0.308536
TEL→CONGDS	3.095316*	0.023507
CONSVS→FIN	0.465648	1.488708
FIN→CONSVS	0.026560	3.246060*
CONSVS→TEL	0.014002	0.038071
TEL→CONSVS	1.825566	0.252877
FIN→TEL	0.051814	1.327027
TEL→FIN	0.191303	1.051570

Panel B: multivariate Granger causality results

	<u>Pre-crisis period</u>	<u>Post-crisis period</u>
TEL, CONSVS, CONGDS & BASMAT → FIN	1.462853	3.736870
FIN, CONSVS, CONGDS & BASMAT → TEL	1.024455	2.548292
FIN, TEL, CONGDS & BASMAT → CONSVS	2.293877	4.168848
FIN, TEL, CONSVS & BASMAT → CONGDS	4.131618	3.467807
FIN, TEL, CONSVS & CONGDS → BASMAT	1.543121	1.651601

Notes: The null hypothesis is rejected based on the χ^2 statistics. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

5.4.1.4 Granger causality

As the Johansen cointegration test suggests that there is no cointegration relationship among the five economic sectors of Greece, we can test the existence of short-run relationships between the returns of the sectors by using an unrestricted VAR model and the

Granger Causality test. As previously mentioned, the number of lags that this VAR model needs in order to ensure that the residuals do not have serial correlation is one. The results of the Granger Causality test are displayed in Table 4. From Panel A of Table 4 we can see that there are four causality relationships which run from BASMAT, FIN and TEL to CONGDS and from FIN to CONSVS at the 10% significance level showing that the CONGDS sector is more likely to be caused by other sectors. However, at our chosen significance level of 5% there are no causality relationships. Finally, as it is shown in Panel B of Table 4 the null hypothesis of no block exogeneity of each sector with respect to the remaining sectors (joint causality) is not rejected at 5% significance level for neither of the two sub-periods.

5.4.1.5 Impulse responses and variance decomposition

The next step is to check the impulse responses of the system of variables by using the generalized impulse responses. The impulse responses examine how a shock to a particular variable at time t , affects the general system of variables in t and future periods. Namely, how fast and towards which direction (positively or negatively) a dependent variable reacts to a shock which happens to another variable. Any shock to a dependent variable should decrease gradually to zero as the underlying variables being modelled are stationary; otherwise there is a permanent effect. The Figure 1 in Appendix F.1 shows the impulse responses between pairs of Greek sectors for the pre-crisis period. The first thing to notice is that the initial reaction of all the variables to a positive shock to another variable is positive and that after one and a half month to two and a half months every shock has died away. Their reaction after the first month differs among different pairs of variables with half of them reacting positively and the rest negatively. Another thing that can be noticed is that all the responses of the TEL, or to the TEL, sector produce the same pattern of positive reaction during the first month and slightly negative during the second month. The Figure 2 in

Appendix F.1 present the generalized impulse responses in the post-crisis period; the results suggest that all the variables react positively from the beginning of the first month and they remain positive until the shocks fade out. However, the responses of the TEL and CONSVS sectors to the FIN sector are different during the pre-crisis period and the rest of the responses of the post-crisis period. Their reaction to a positive shock to another sector is positive during the first month and negative during the second month until the shock equals zero. Moreover, the time needed for the shock to fade out is the same as during the pre-crisis period of approximately one and a half to two and a half months.

In addition to the impulse responses we also perform forecast error variance decomposition. Variance decomposition investigates the proportion of a variable's variance movements which are caused by its own innovations against the proportion of variance movements which are due to shocks in other variables of the system. Based on this proportion, each variable of the system can be classified as exogenous or endogenous: that is, a variable is called exogenous when its variance decomposition is mainly due to its own shocks and endogenous when its variance decomposition is more than 50% explained by shocks in other variables. As for the variance decomposition we use the Cholesky factorisation to orthogonalise the VAR innovations in order that the shocks are uncorrelated contemporaneously, we need to pay attention to the ordering of the sectors in the VAR model; that is, this method can provide very different results depending on the ordering of the variables in the VAR model. To determine the ordering we start by computing the average correlation of each sector with the others to find the level of exogeneity of each one. The ordering is then set from the sector with the lowest average correlation to the sector with the highest average correlation. The pre-crisis period's results (Appendix F.2) indicate that the TEL and CONGDS sectors are the most exogenous variables out of all. This means that the shocks that happen to these sectors are consequences of innovations in themselves. Over a period of five months CONSVS, BASMAT and FIN are endogenous sectors and less

than 50% of their variances are explained by innovations that happen in these sectors. For instance, the variance of CONSVS is 31% and 18% explained by innovations in the CONGDS sector and the TEL sector respectively while 49% is explained by innovations in itself. On the contrary, during the post-crisis period the variables that are considered as exogenous are the CONGDS, TEL and FIN and the rest are considered as endogenous because their variance depends mostly on the other sectors (Appendix F.2). To sum up, the CONGDS and TEL sectors are exogenous through the whole period tested, the BASMAT and CONSVS are endogenous and the FIN sector changes between pre-crisis and post-crisis period.

5.4.1.6 Time-varying volatilities and correlations

Appendix F.3 shows the time-varying correlation and time-varying volatility results for the whole period tested. From the time-varying correlation it is noticeable that we can divide these results into two groups. The first group includes the correlation between TEL and CONGDS, TEL and CONSVS, and TEL and BASMAT. The second group includes the correlations of the rest of the pairs. The correlations between the pairs of the first group illustrate that the correlations of TEL with CONGDS, CONSVS and BASMAT were significantly lower during the period of 1999 to 2001 but they increased a lot after 2001 and they are still at high levels. The correlations of the pairs in the second group were much higher during the pre-crisis period but they decreased considerably by an average of 11% during the post-crisis period and they are still gradually diminishing. Finally, a very important result of the time-varying correlations is that the correlations between all the different pairs of variables are positive through the whole period of 1998 to 2016.

The results of time-varying volatility suggest that the FIN, TEL and CONGDS sectors are the ones that experienced high and sudden increases in the volatility of their returns, while the CONSVS and BASMAT did not have significant surges in their volatility. Specifically, CONGDS

and CONSVS experienced high levels of volatility during the beginning of the pre-crisis period and they remain at low levels since then. Conversely, FIN and TEL had extreme rises in their volatility during the post-crisis period. This sharp increase in the volatility of the FIN sector is a result of the financial crisis in Europe which began on 2008 as already mentioned. The FIN sector reached its highest peak of 0.55 at the middle of 2015 probably due to the referendum, which took place in June of 2015, about the austerity measures and the possibility of Greece rejecting a further bailout loan. In addition to that, the enforcement of capital controls on the residents by the government made the economic market even more uncertain and volatile.

5.4.2 Italy

5.4.2.1 Descriptive statistics

Table 5 provides a summary of the descriptive statistics of the five most important sectors in Italy. These are the consumer goods (CONGDS), the financials (FIN), the oil and gas (OG), the telecommunications (TEL) and the utilities (UTI) sectors. Panel A contains the statistical characteristics for the whole period tested and the results show that four out of the five sectors have a negative average return, while only the CONGDS sector has a very low positive return of 0.11%. As expected, the negative returns are accompanied by high standard deviation (risk). Moreover, the skewness of all the sectoral data is negative which implies that there is higher possibility for negative returns to exist, rather than positive. In addition, the values of skewness and kurtosis confirm the results of the Jarque-Bera test, that the data are not normally distributed. That is the prices of the sectoral data are not distributed approximately 50% above and 50% below the mean value.

Table 5

Descriptive statistics of returns in different periods – Italy

Period	CONGDS	FIN	OG	TEL	UTI
<i>Panel A Whole Period 01/06/1998 – 01/11/2016</i>					
Mean	0.0011	-0.0050	-0.0003	-0.0056	-0.0002
Median	0.0097	0.0028	0.0064	-0.0031	0.0048
Maximum	0.2332	0.2149	0.1619	0.2510	0.2400
Minimum	-0.2617	-0.2588	-0.2155	-0.2756	-0.1792
Std. deviation	0.0709	0.0799	0.0590	0.0820	0.0535
Skewness	-0.6578	-0.5883	-0.4372	-0.1732	-0.1814
Kurtosis	4.6401	3.9406	3.3895	3.6847	5.1275
Jarque-Bera	40.7083***	20.8962***	8.4409**	5.4221*	42.893***
<i>Panel B Pre-crisis Period 01/06/1998 – 01/08/2007</i>					
Mean	-0.0005	0.0021	0.0065	-0.0025	0.0026
Median	0.0111	0.0084	0.0186	-0.0022	0.0055
Maximum	0.1432	0.1650	0.1164	0.2510	0.2400
Minimum	-0.2617	-0.2257	-0.2155	-0.2153	-0.1792
Std. deviation	0.0709	0.0635	0.0542	0.0818	0.0538
Skewness	-1.1455	-1.1015	-0.8784	0.1154	0.0381
Kurtosis	5.1329	5.4986	4.3974	3.9140	7.1006
Jarque-Bera	44.9093***	50.8568***	23.0974***	4.0734	77.0957***
<i>Panel C Post-crisis Period 01/09/2007 – 01/11/2016</i>					
Mean	0.0030	-0.0124	-0.0071	-0.0083	-0.0032
Median	0.0100	-0.0062	-0.0048	-0.0057	0.0003
Maximum	0.2332	0.2149	0.1619	0.1659	0.1148
Minimum	-0.2160	-0.2588	-0.1783	-0.2756	-0.1416
Std. deviation	0.0716	0.0935	0.0634	0.0826	0.0534
Skewness	-0.1906	-0.2664	-0.0788	-0.4355	-0.3981
Kurtosis	4.0483	3.0596	2.9248	3.4037	2.9901
Jarque-Bera	5.7026	1.3176	0.1397	4.2241	2.9059

Notes: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. The summary statistics represent the five most important market economic sectors of Italy: Consumer Goods (CONGDS), Financials (FIN), Oil & Gas (OG), Telecommunications (TEL) and Utilities (UTI). The null hypothesis of the Jarque-Bera test is the skewness to be equal to zero and the kurtosis to be equal to three.

By comparing the data of Panels B and C in Table 5 one can notice that the mean return of the sectors have almost the exact opposite pattern comparing the pre-crisis and post-crisis periods. During the pre-crisis period the mean of the CONGDS and the TEL sectors were negative and the rest of the sectors were positive, while during the post-crisis period only CONGDS had a positive mean return and the rest of the sectors had a negative one. The sector which had the largest decrease in its return was the FIN sector by 694%. This is a sign

of how much the FIN sector was affected by the global financial crisis. The rest of the sectors, except the CONGDS sector which noted an increase between the pre and post-crisis period, were reduced by nearly 221%. The risk of the sectors remains almost the same between the two different periods tested. The only exception is in the FIN sector where the risk increased significantly more than in the other sectors. This is not surprising as it also experienced the largest decrease in its mean return. Furthermore, the negative skewness results indicate that there were more possibilities to experience high negative returns and the positive skewness shows the exact opposite of larger possibilities to have large positive returns. The skewness of the TEL and UTI sectors is positive during the pre-crisis period, while it is negative for all the sectors during the post-crisis period. This is probably a result of the general financial crisis which usually gives rise to negative returns. In addition, we observe from the results of the Jarque-Bera normality test that the sectoral returns of Italy are not normally distributed in the pre-crisis period (except for the TEL sector), but they are all follow a normal distribution after the global financial crisis of 2008. That means that these returns in the post-crisis period were distributed evenly around the average return.

5.4.2.2 Dickey – Fuller GLS unit root test

The first step in the analysis of the Italian sectors is to check them for stationarity. In order for the possibility of a long-run equilibrium to exist among the data, it is necessary for the level prices of the sectors to be non-stationary. After checking the graph of each individual sector of both sub-periods for the possibility of trend existence, we see that only the OG sector exhibits a trend during the pre-crisis period, while the CONGDS and OG sectors exhibit a trend during the post-crisis period. Hence we include trend in the Dickey – Fuller GLS regressions for these sectors. The DF-GLS unit root test results for both level prices and first differences are displayed in Table 6. Panels A and B contain the results for the pre-crisis and

post-crisis periods respectively. Based on these results we find that none of the data series are stationary in level values for both the pre-crisis and post-crisis periods. All of them are integrated of order one (I(1)), which means that their levels have a unit root, but they are stationary in first differences. This gives rise to the possibility of the existence of a long-run equilibrium relationship between the variables.

Table 6
Dickey-Fuller GLS unit root test – Italy

Period		CONGDS	FIN	OG	TEL	UTI
<i>Panel A – Pre-crisis period</i>						
<i>Log prices</i>						
	Lag	0	0	0	0	0
	τ_{μ}	-0.9191	-1.3002	-1.7724	-1.3758	-1.2089
<i>Log returns</i>						
	Lag	0	1	1	0	2
	τ_{μ}	-9.2954***	-7.0223***	-	-	-2.2002**
				6.3924***	8.8035***	
<i>Panel B – Post-crisis period</i>						
<i>Log prices</i>						
	Lag	1	0	0	0	0
	τ_{μ}	-1.4908	-0.0632	-0.1798	-0.4203	-0.7334
<i>Log returns</i>						
	Lag	0	0	0	0	0
	τ_{μ}	-7.8238***	-9.2789***	-	-	-
				8.8158***	9.2419***	8.0065***

Notes: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. The shaded cells indicate trend existence. The critical values for Dickey-Fuller GLS unit root test when there is no trend are -2.586350, -1.943796 and -1.614784 at 1%, 5% and 10% levels of significance, respectively. In the case of trend existence, the critical values change to -3.568000, -3.020000 and -2.730000, respectively.

5.4.2.3 Johansen cointegration test

As the sectoral data are I(1) for both subsamples, they can be tested for cointegration relationships by applying the Johansen cointegration test to the levels of the prices. The number of lags for the test is selected based on an unrestricted VAR which uses the level prices and it is equal to the number for which there is no serial correlation left in the residuals. For the first sub-sample, five lags were used. The Akaike information criterion (AIC) recommended one lag. However, one lag is not sufficient to remove the serial correlation

from the residuals. After checking other number of lags, we end up using five lags. For the second sub-period the serial correlation is eliminated for two lags. The outcome of the Johansen Trace test for cointegration is displayed in Table 7 and it suggests that there is a cointegration relationship among the series during both the pre-crisis and post-crisis periods. This means that there is a long run equilibrium relationship among the five sectors of Italy over both sub-periods. The speed of adjustment of the sectoral variables towards this equilibrium after a shock to their price can be produced by the vector error correction model (VECM). This should be negative and significant in order for the variables to return back to their equilibrium. In the opposite case that it is positive and after an exogenous shock, the variables will move continuously away from the equilibrium. From the VECM model, we can test the short-run relationships of the sectors for each period by applying the Granger causality test.

Table 7
Multivariate Johansen cointegration test – Italy

	Pre-crisis		Post-crisis	
	λ_{trace}	λ_{max}	λ_{trace}	λ_{max}
$r = 0$	78.4828**	42.8660**	73.7352**	31.3476
$r \leq 1$	35.6168	15.7990	42.3876	25.8014
$r \leq 2$	19.8177	12.7665	16.5862	9.8056
$r \leq 3$	7.0513	3.7048	6.7806	6.6692
$r \leq 4$	3.3464	3.3464	0.1113	0.1113

Notes: The optimal lag length for the cointegration testing is defined by the Schwarz information criterion (SC). ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

5.4.2.4 Granger causality

Although the Johansen cointegration test suggests that there is a long-run equilibrium relationship among the sectors, it does not give any information about the possibility of short-run relationships between them. The Granger causality test is used on the logarithmic

returns of the sectors and it shows that there are several causality relationships among them during the pre-crisis period. Specifically, from panel A of Table 8 one can notice that there is causality running from the UTI sector to the CONGDS sector at the 1% significance level (given that the test estimator is greater than the critical value at the 1% significance level). This means that the past values of UTI sector contain information which can be important for the prediction of the future stock price movements of the CONGDS sector. Similarly, it is found that the UTI sector is Granger-caused by both the FIN and TEL sectors at the 5% significance level, but not the other way around. Furthermore, there are bi-directional causal relationships between the OG & TEL sectors and between the TEL & CONGDS sectors at the 10% and the 5% significance level, respectively. Although, the sectors have a long run equilibrium during both sub-periods and there are many short-run relationships during the pre-crisis-period, there are only two causal relationships during the post-crisis period. This is the causality which runs from the OG sector to the TEL sector at the 5% significance level and from the CONGDS sector to the FIN sector at the 10% significance level. It can be noticed that almost all the pairwise causal relationships which exist before the financial crisis no longer exist after it. Finally, the Granger causality test examines the probability of the lagged values of one sector to improve the forecast of another sector but also the probability of the lagged values of four out of the five sectors to jointly improve the forecast of the fifth one. Panel B of Table 8 gives the results of the multivariate Granger causality test for both the pre-crisis and the post-crisis periods. In particular, there are four strong jointly causal relationships among the sectors in the period before the financial crisis. For instance, the lagged returns of the FIN, OG, TEL and UTI sectors jointly cause the CONGDS sector at the 1% significance level. That means that the FIN, OG, TEL and UTI sectors carry past information which can be used to forecast the future values of the CONGDS sector. Similarly, the FIN sector, the TEL sector and the UTI sector are caused jointly by the remaining variables. The only exception is the case of OG sectors, where there is no jointly causal relationship running from the

CONGDS, FIN, TEL and UTI sectors to the OG sector. On the contrary, after the financial crisis there is no jointly causal relationship for none of the Italian sectors.

Table 8
Granger causality test results – Italy

	<u>Pre-crisis period</u>	<u>Post-crisis period</u>
OG→CONGDS	4.2142	0.2687
CONGDS→OG	1.1185	1.3512
OG→UTI	1.0490	0.1197
UTI→OG	7.2673	0.0024
OG→FIN	1.6469	0.0535
FIN→OG	7.3341	1.0529
OG→TEL	8.9472*	4.8244**
TEL→OG	8.9935*	0.4311
CONGDS→UTI	7.2265	0.4933
UTI→CONGDS	14.4828***	0.2239
CONGDS→FIN	6.8791	2.7814*
FIN→CONGDS	7.5850	0.0491
CONGDS→TEL	10.7912**	0.2017
TEL→CONGDS	23.9916***	0.3596
UTI→FIN	6.0045	0.0659
FIN→UTI	10.7787**	0.0003
UTI→TEL	4.2254	0.1070
TEL→UTI	12.7167**	0.0849
FIN→TEL	6.4176	1.1401
TEL→FIN	7.8641	0.0029
<i>Panel B: multivariate Granger causality results</i>		
	<u>Pre-crisis period</u>	<u>Post-crisis period</u>
TEL, UTI, CONGDS, FIN → OG	20.2268	5.2402
OG, UTI, CONGDS, FIN → TEL	44.0651***	6.3553
OG, TEL, CONGDS, FIN → UTI	52.5870***	1.0735
OG, TEL, UTI, FIN → CONGDS	37.9639***	1.7344
OG, TEL, UTI, CONGDS → FIN	32.0502***	2.8500

Notes: The null hypothesis is rejected based on the χ^2 statistics. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

5.4.2.5 Impulse responses and variance decomposition

Following the Granger causality test we now examine the impulse responses of the system and the variance decomposition of the sectoral returns. As mentioned earlier, the impulse responses used in this study are the generalised ones and they show how a sector of the system reacts to a shock which occurs in another sector of the system independently of the variable ordering. For the variance decomposition which illustrates the proportion of movements in a sector's price due to its own innovations, also called the "degree of exogeneity", a Cholesky decomposition is performed on the covariance matrix of the shocks and therefore these results depend on the ordering of the variables in the system.

Figure 1 in Appendix G.1 shows the graphs of the impulse responses for the pre-crisis period and one of the first things to notice is that all of the sectors have an initial positive reaction to shocks in other sectors. However, some of them have a negative response after the first month. These are the response of the TEL sector to the OG and UTI sectors, and the response of the UTI sector to the OG and TEL sectors. The rest, which are the majority of the sectors, have a positive response to a shock which occurs to other sectors. One issue which is really interesting is whether a shock has a permanent effect on another sector or whether it is transitory. This is a question which arises from the fact that the sectors have been found to be cointegrated during this period of time. In the case that they were not cointegrated the shock effect would be temporary and it would gradually approach zero. Out of all the responses only the response of the TEL sector to the CONGDS and FIN sectors, and the response of the UTI sector to the CONGDS and FIN sectors are found to have a transitory effect of a shock on another sector. Finally, for almost all the responses at least 10 months are needed for the shock's effects to die out or to stabilise. The results of the post-crisis period are exhibited in Figure 2 in Appendix G.1 where one can immediately notice that all the responses are positive with the majority of them being stabilised after a range of five to

seven months. All the sectors react in a more discernible way than during the pre-crisis period and the effects of other sectors on them are permanent, meaning they do not die out. The reactions of the OG, TEL, FIN and CONGDS sectors are slightly more positive and there are less fluctuations to their responses until the shock effects are stabilised. Furthermore, the reactions of the UTI sector to a positive shock to the rest of the sectors are positive while they are negative before the financial crisis. An interesting outcome of the comparison between the pre-crisis and post-crisis period is that the effect of a shock dies out much faster during the post-crisis period.

The tables in Appendix G.2 demonstrate the results of the variance decomposition of the sectoral data for the pre-crisis and post-crisis periods, respectively. During the period before the financial crisis the most exogenous sector was the OG sector and the most endogenous was the FIN sector, while after the financial crisis the most endogenous remained the FIN sector but the most exogenous was the TEL sector. Also during the pre-crisis period OG, TEL and UTI are considered as exogenous and the rest of the sectors as endogenous. For example, 69% of the innovations of the UTI sector approximately depend on shocks in the UTI sector and 31% of them are due to shock in other sectors. During the post-crisis period there are some changes. Firstly, all the sectors seem to be a bit more exogenous. That is, their forecast error variance is mostly due to innovations to themselves. The only exception is the FIN sector which appears to be more endogenous compared to the pre-crisis period with only 25% of its forecast error variance being explained by innovations to the same sector. In addition, the FIN sector seems to be affected a lot by shocks in the TEL and CONGDS sectors. The influence of TEL to FIN sector seems to decrease gradually while the CONGDS sector appears to have a growing influence on the FIN sector, which is also supported by the causality relationship which runs from CONGDS to the FIN sector.

5.4.2.6 Time-varying volatilities and correlations

In the left column of the figure in Appendix G.3 one can see the graphs of the time-varying volatility of the Italian sectors. The results imply that there were not large fluctuations in the volatility of the sectors either before or after the financial crisis of 2008. There is only a slight increase in the volatility after 2008, especially for the FIN, TEL and OG sectors.

The middle and right columns of the Figure in Appendix G.3 summarise the results of the pairwise Dynamic Conditional Correlation (DCC) for the five Italian sectors. To begin with, there are only four pairs which experienced an increase in their correlation after the financial crisis of 2008. In particular, the average correlations between the CONGDS and UTI sectors, OG and UTI sectors, TEL and UTI sectors and UTI and FIN sectors were increased after 2008, while the average correlation of the rest of the pairs tested remained almost the same as during the period before 2008. Finally, we notice that during the first three years after the start of the financial crisis of 2008, there is either a visible increase or decrease in all the pairs. All of these movements are positive, except for the correlation between TEL and FIN, TEL and UTI, and TEL and OG. Generally, the correlation of the TEL sector with the rest of the sectors seems to have been reduced immediately after the financial crisis of 2008.

5.4.3 Portugal

5.4.3.1 Descriptive statistics

Table 9 summarises the statistical properties of the sectors' returns defined as the first differences of their logarithmic prices. Panel A of Table 9 contains the results for the whole period tested. The CONSVS sector is the only one which exhibits a positive but also small return with a relatively low risk level (standard deviation), while the other four sectors are characterised by negative returns accompanied by high standard deviations something

which is expected as the average returns are negative and this results in a higher degree of uncertainty and risk.

Table 9

Descriptive statistics of returns in different periods – Portugal

Period	CONSVS	FIN	IND	TEL	UTI
<i>Panel A Whole Period 01/06/1998 – 01/11/2016</i>					
Mean	0.0003	-0.0151	-0.0042	-0.0098	-0.0020
Median	0.0015	-0.0052	-0.0001	-0.0024	-0.0040
Maximum	0.1485	0.2983	0.1657	0.4179	0.2244
Minimum	-0.2730	-0.4912	-0.3278	-0.3839	-0.2415
Std. deviation	0.0624	0.0996	0.0691	0.0951	0.0642
Skewness	-0.6767	-0.5882	-0.9849	-0.3839	-0.1260
Kurtosis	4.9431	5.8200	6.3515	6.4391	4.3616
Jarque-Bera	51.631***	85.974***	139.16***	114.34***	17.656***
<i>Panel B Pre-crisis Period 01/06/1998 – 01/08/2007</i>					
Mean	-0.0002	0.0033	0.0063	0.0002	-0.0010
Median	0.0048	0.0066	0.0059	-0.0019	-0.0079
Maximum	0.0944	0.2320	0.1380	0.4179	0.2244
Minimum	-0.1821	-0.2306	-0.2256	-0.3644	-0.1261
Std. deviation	0.0494	0.0560	0.0559	0.0978	0.0625
Skewness	-0.5670	-0.0818	-0.5810	0.2003	0.6154
Kurtosis	3.6151	7.4712	5.1564	6.8939	4.1851
Jarque-Bera	7.6283**	91.752***	27.503***	70.230***	13.380***
<i>Panel C Post-crisis Period 01/09/2007 – 01/11/2016</i>					
Mean	0.0012	-0.0332	-0.0144	-0.0196	-0.0028
Median	0.0006	-0.0450	-0.0082	-0.0030	0.0005
Maximum	0.1485	0.2983	0.1657	0.1532	0.1387
Minimum	-0.2730	-0.4912	-0.3278	-0.3839	-0.2415
Std. deviation	0.0734	0.1273	0.0792	0.0922	0.0663
Skewness	-0.6902	-0.2191	-0.9268	-1.1362	-0.7471
Kurtosis	4.4971	3.8748	5.6362	5.3332	4.3771
Jarque-Bera	19.007***	4.3879	47.599***	48.617***	18.926***

Notes: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. The summary statistics represent the five most important market economic sectors of Portugal: Consumer Services (CONSVS), Financials (FIN), Industrials (IND), Telecommunications (TEL) and Utilities (UTI). The null hypothesis of the Jarque-Bera test is the skewness to be equal to zero and the kurtosis to be equal to three.

The FIN sector has the most negative average return of 1.5097% and it also has the higher risk. All returns have negative skewness which indicates that the distribution has a long right tail, meaning that there are more negative returns than positive. The kurtosis values are higher than three implying that the distribution has longer and fatter tails than the normal

distribution. The values of skewness and kurtosis confirm the results of the Jarque-Bera normality test which examines if the returns follow a normal distribution. As its null hypothesis of normality is rejected, it is implied that the returns are not normally distributed. Panels B and C of the same table provide statistical details of the sectors for the pre-crisis and post-crisis period, respectively. Before the financial crisis three out of the five sectors have positive mean returns; these are the FIN, IND and TEL with IND having the highest mean return and the second lowest risk. The other two sectors, CONSVS and UTI, have negative average rate of returns, but not the highest risk. Out of the sectors with a positive mean return the IND sector has the highest return and the lowest risk – a result which goes against the theory that higher return is accompanied by higher risk. Furthermore, the skewness of the data is negative for CONSVS, FIN and IND indicating the high possibility of negative returns; the skewness of the other two sectors is positive. In addition, for all the sectors kurtosis is larger than three but it is very close to it in the case of CONSVS. Thus, we can say that the CONSVS sector has a mesokurtic distribution and the rest a leptokurtic distribution. Again the values of skewness and kurtosis ensure the result of the Jarque-Bera test that the returns do not follow a normal distribution. As mentioned above, panel C shows the results of the post-crisis period. To begin with, the majority of the sectors have a negative mean return with only exception the CONSVS sector. The latter seems to have a significant increase in its returns because, although during the pre-crisis period it has a negative average return, over the whole period it is characterised by a positive return. Thus, its returns after the financial crisis are higher than the returns it experienced during the period before the crisis. The rest of the sectors have a negative average return indicating the effect of the financial crisis on their prices. The sector which experienced the highest decrease in its average return is the TEL followed by the FIN where both have their mean returns reduced by more than 1000%. Skewness is negative for all the sectors and kurtosis is close but not equal to three. Thus, they do not follow a normal distribution except the FIN sector which seems to follow

such a distribution as the null hypothesis of normality by the Jarque-Bera test is not rejected at the 1% significance level.

5.4.3.2 Dickey-Fuller GLS unit root test

Table 10 shows the results of the Dickey-Fuller unit root test. After checking the graph of each individual sector in order to inspect the possible existence of a trend in the data, the results show that all the sectors are characterised by a trend apart from the TEL sector in the pre-crisis period and the UTI sector in the post-crisis period. Testing the data for the Portuguese sectors for stationarity with the Dickey-Fuller unit root test for both the pre-crisis and post-crisis periods, we conclude that all of them, except the TEL sector, are integrated of order one as the null hypothesis of unit root is not rejected when the test is applied to the levels of the prices but is rejected when it is applied to their first differences.

5.4.3.3 Johansen cointegration test

Since based on the Dickey-Fuller GLS unit root test all the data series are found to be integrated of the same order except TEL on the pre-crisis period, the next step in the analysis is to test the sectors for cointegration. The Johansen cointegration results for the pre and post crisis periods are provided in Table 11 below. For the pre-crisis period the Johansen cointegration test is used with all the sectors included except TEL, as the variables included in the test should be integrated of the same order. The number of lags used for the Johansen cointegration test is 1 (selected by the AIC) which appears to be sufficient in terms of no remaining serial correlations. Both the Trace and Maximum Eigenvalue statistics suggest that there is no cointegration among the data (long-run equilibrium) as the first null hypothesis

of no cointegration ($r = 0$) is not rejected at 5% significance level by neither of the tests. That is, the test statistics were smaller than the critical values of the test.

The Johansen test is also used for the post-crisis period. This time all the sectors are included as all of them are I(1). The serial correlation is eliminated with three lags which is not consistent with the AIC recommendation. The AIC information criterion suggests one lag, but this is not enough for the serial correlation to be eliminated. After trying two lags, which were not enough either, we conclude that the number of lags we need is three. Based on the Trace test of Johansen (1991) there is a long-run equilibrium relationship among the five main Portuguese sectors for the post-crisis period. To sum up, the sectors are not cointegrated during the pre-crisis period but they are during the post-crisis period.

Table 10
Dickey-Fuller GLS unit root test – Portugal

Period		CONSVS	FIN	IND	TEL	UTI
<i>Panel A – Pre-crisis period</i>						
<i>Log prices</i>						
	Lag	1	1	0	0	0
	τ_μ	-0.4642	-0.6378	-0.0406	-2.2531**	-0.3376
<i>Log returns</i>						
	Lag	0	0	0		0
	τ_μ	-8.3888***	-8.5859***	-	-	-
				9.7977***	10.525***	5.4417***
<i>Panel B – Post-crisis period</i>						
<i>Log prices</i>						
	Lag	0	1	0	0	3
	τ_μ	-1.6414	-2.2092	-2.2570	-1.3618	-1.1132
<i>Log returns</i>						
	Lag	0	0	0	0	0
	τ_μ	-9.4287***	-8.3708***	-	-	-
				9.4615***	10.491***	11.272***

Notes: ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. The shaded cells indicate trend existence. The critical values for Dickey-Fuller GLS unit root test when there is no trend are -2.586350, -1.943796 and -1.614784 at 1%, 5% and 10% levels of significance, respectively. In the case of trend existence, the critical values change to -3.568000, -3.020000 and -2.730000, respectively.

Table 11
Multivariate Johansen cointegration test - Portugal

	Pre-crisis		Post-crisis	
	λ_{trace}	λ_{max}	λ_{trace}	λ_{max}
$r = 0$	42.3115	20.6397	74.3143**	31.98719
$r \leq 1$	21.6718	15.9042	42.3272	19.46959
$r \leq 2$	5.7676	5.7105	22.8576	17.39109
$r \leq 3$	0.0572	0.0572	5.4664	5.059822

Notes: The optimal lag length for the cointegration testing is defined by the Akaike information criterion (AIC). ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

5.4.3.4 Granger causality

As a consequence of the missing cointegration among the data on the pre-crisis period, we will use an unrestricted VAR model with the first differences to examine the data for short-run relationships (causality) during this period. We also use a VECM model with the level prices on the post-crisis period as we found cointegration. There are some Granger causality relationships among the data. First, during the pre-crisis period the CONSVS sector seems to Granger-cause the UTI sector as the null hypothesis of CONSVS to not cause the UTI sector is rejected at the 5% significance level. Similarly, the FIN sector causes the IND sector at the 10% significance level. Moreover, as it can be seen from panel B of Table 12, for the pre-crisis period the null hypothesis of no Granger causality to each sector from the remaining sectors is not rejected for any of the five sectors, indicating that there are no joint causality relationships among the data.

Secondly, during the post-crisis period the short-run relationships have doubled. The causal relationship of the FIN sector to the IND sector still exists at the 10% significance level. There are also some new short-run relationships. One unidirectional from the TEL to the CONSVS sector and a bi-directional one between the TEL and the UTI sectors. Thus after the financial

crisis the causality relationships between the data have risen and there are new relationships created mostly from and to the TEL sector.

Table 12
Granger causality test results - Portugal

Panel A: Pairwise Granger causality results

	<u>Pre-crisis period</u>	<u>Post-crisis period</u>
IND→CONSVS	0.6783	0.7306
CONSVS→IND	0.0110	1.6508
IND→UTI	0.3786	0.5602
UTI→IND	0.0202	2.3478
IND→FIN	0.3939	0.3942
FIN→IND	3.1025*	2.7744*
IND→TEL	0.6858	0.0154
TEL→IND	0.1028	0.2728
CONSVS→UTI	4.1099**	1.7659
UTI→CONSVS	1.2316	0.3029
CONSVS→FIN	2.1118	0.6434
FIN→CONSVS	2.5021	1.5831
CONSVS→TEL	0.5494	0.1641
TEL→CONSVS	1.8096	3.2550*
UTI→FIN	1.0830	0.0233
FIN→UTI	0.0004	0.4086
UTI→TEL	0.1559	4.1252**
TEL→UTI	0.6578	4.0128**
FIN→TEL	0.0319	1.1540
TEL→FIN	0.3521	0.5951

Panel B: multivariate Granger causality results

	<u>Pre-crisis period</u>	<u>Post-crisis period</u>
FIN, IND, TEL & UTI → CONSVS	5.3518	4.2873
CONSVS, IND, TEL & UTI → FIN	5.4376	2.3113
CONSVS, FIN, TEL & UTI → IND	5.4654	14.1456***
CONSVS, FIN, IND, & UTI → TEL	1.1650	5.0167
CONSVS, FIN, IND & TEL → UTI	4.3286	5.2347

Notes: The null hypothesis is rejected based on the χ^2 statistics. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Lastly, the combined Granger causality from four sectors to the fifth one is rejected in all cases apart from the IND sector which seems to be caused by the other sectors at the 1%

significance level. That is, the past values of the CONSVS, FIN, TEL and UTI sectors can be helpful in the prediction of the future values of the IND sector.

5.4.3.5 Impulse responses and variance decomposition

The impulse responses for the pre-crisis and post-crisis periods are displayed respectively in Figure 1 and Figure 2 of Appendix H.1. The first results of the impulse response analysis indicate that all the sectors before and after the financial crisis react positively during the first month after an exogenous shock to another sector of the system. The majority of the sectors' responses during the pre-crisis period remain positive until they reach zero. However, there is the case of the UTI sector responses to the IND, TEL and FIN sectors, as well as the case of the TEL sector responses to UTI and FIN which have a positive reaction in the first month but after the first month their reactions become negative until the shock's effects die out. The time needed until the responses reach zero is between 2.5 and 3.5 months. By comparing the results of the post-crisis period to the results of the pre-crisis period it is observed that during the post-crisis period all the responses of the Portuguese sectors to random shocks to the rest of the sectors in the system are positive and the shocks' effects are permanent for all the sectors. Furthermore, the average period needed so that the shocks effect to a sector is stabilised is highly increased as it lies between 11 and 15 months.

In addition to the results of the impulse responses, we have the variance decomposition of the sectors. Appendix H.2 provides the results of the variance decomposition analysis of each sector for both sub-periods tested. During the pre-crisis period three out of the five sectors seem to be exogenous. These are the IND, CONSVS and UTI which explain their variance decomposition by their own innovations at a degree of 88%, 75% and 62%, respectively. The rest are found to be endogenous sectors as at the 5-month horizon 52% of the TEL sector

and 58% of the FIN sector innovations are explained by the remaining sectors, with the CONSVS sector being the most important in their variance explanation. By analysing the results of the variance decomposition for the period after the financial crisis, we notice that the level of exogeneity has increased for all the sectors, except the CONSVS sector. This means that the innovations of each sector are mostly due to their own shocks. Moreover, the IND sector remains the most exogenous and the FIN the most endogenous, the same as for the pre-crisis period.

5.4.3.6 Time-varying volatilities and correlations

The graph in Appendix H.3 represents the results of the time-varying volatility and the pairwise time-varying correlation of the sectors. With regard to the results of time-varying volatility, a significant increase in the volatility of all sectors is highlighted after 2008 possibly due to the financial crisis that influenced Portugal during this period. Moreover, the FIN sector has the highest risk with a volatility as high as 0.15, while the rest of the sectors have a volatility of less than 0.02. This is a sign of how risky the FIN sector is considered and how volatile its value was after 2008. The next most risky sector is CONSVS, the volatility of which has the highest peak of 0.018 in 2008 but after this year its value is not so volatile. To conclude, the Portuguese sectors are not characterised by high levels of volatility with the FIN sector being the most volatile one.

The final step in the analysis is to check the time-varying correlations of the sectors. The results in the two right-hand columns of the Figure in Appendix E.3 show that the correlations between the sectors are not very volatile. However, there are some peaks in 2008, 2012 and 2014. In 2008, the correlation increased for all the different sector pairs except the FIN sector with the IND and the TEL sector, and the UTI and TEL sectors. In 2012, the correlation between the IND sector and the remaining sectors decreased significantly, while the

correlation of the UTI and TEL sectors surged by 0.15 units. Finally, in 2015 the correlation of the FIN sector with the TEL, UTI and CONSVS sectors, as well as that between the TEL and CONSVS sectors increased. In a nutshell the correlations only exhibit a moderate degree of fluctuation but there are some brief periods during the period examined where they increased significantly.

5.5 Summary and conclusions

The purpose of this study was to investigate the interdependency relationships among the stock market sector indices of Greece, Italy and Portugal. The data cover the period June 1998 through November 2016 for the five most important market sectors (based on their share in each country's total market capitalisation). The methodology applied includes unit root tests, long-run and short-run relationship tests, impulse response and variance decomposition analyses, and time-varying volatilities and correlations checks.

Our findings indicate that the Greek sectors were influenced greatly after the beginning of the financial crisis since all the mean returns of all the sectors reduced significantly, with the Financials sector being the most influenced – with a loss in its value greater than 1000%. All the sectors are found to be integrated of order one and they do not share a long-run equilibrium relationship for neither of the sub-periods under investigation. However, there are some short-run relationships which show that the Consumer Services sector has the highest possibility to be caused by the other sectors.

Moreover, the Italian sectors show that most of their average returns are negative, especially after the financial crisis of 2008. The FIN sector has again the highest decrease out of all the Italian sectors in its average returns but also the highest risk. Only in the case of Italy there is cointegration found among the data for the pre-crisis period which is accompanied by

many causality relationships. These are mainly pairwise relationships which run from the Financials and Oil & Gas sectors to the Telecommunication sector and from Telecommunication and Utilities sectors to the Consumer Goods sector. Moreover, there are short-run dynamics of joint causality during the pre-crisis period running to each individual sector from the remaining ones, except the Oil & Gas sector. However, after the financial crisis there is cointegration but there is only one short-run relationship from the Consumer Goods sector to the Financials sector. This is a sign that when sectors are cointegrated, there is higher possibility for short-run relationships to exist.

Finally, the Portuguese sectors have mostly positive returns before the crisis and negative after it. In addition, the Portuguese Financials sector is the one that was most affected by the financial crisis, as well as in Greece and Italy. The risk of the sectors also significantly increased after the start of the financial crisis at the end of 2007, especially for the Financials sector. The Johansen cointegration test showed that there is no long-run equilibrium among the sectors during the pre-crisis period but there is during the post-crisis period. Moreover, there is a unidirectional causality relationship from the Consumer Services sector to Utilities sector during the first sub-period and a bi-directional causality relationship between the Utilities sector and the Telecommunication sector during the post-crisis period. Moreover, after the financial crisis there is joint causality relationship running to the Utilities sector from the other sectors.

To sum up, we conclude with some very interesting observations. First, the Financials sector is the most sensitive one in these countries, having a substantial decrease in its average return. Second, the Financials and Telecommunication sectors have a significant share of the total countries' market capitalization as they are included in the main sectors of the countries under investigation. Third, the sectors are not strongly cointegrated most of the time and they have limited causality relationships. This can provide diversification benefits to

investors, as there are very weak interdependency relationships (Wang *et al*, 2005). Furthermore, the impulse response functions indicate that in order for the effects of a shock to a sector to die out, an average of two and a half months is needed. The variance decomposition analysis shows that most of the sectors are exogenous rather than endogenous and they experience a much higher volatility after a financial crisis. This is not surprising as when a financial crisis occurs, the degree of uncertainty is significantly increased. Last but not least, the time-varying correlations show that most of the pairwise correlations for Greek sectors have considerably declined over the second sub-period under investigation (after the end of 2007). On the contrary, in the case of the Italian sectors the pairwise correlations have increased, while the correlations between the Portuguese sectors seem to not have changed.

We conclude that there are no interdependency relationships among the sectors of Greece, Italy and Portugal. They are mainly affected in their volatility levels and their pairwise correlations. In addition, we conclude that the volatility levels of the sectors are increased significantly after a financial crisis as the uncertainty is also surged. However, it is not clear how the correlations react after a financial crisis as the results of the time-varying correlations for the countries we examined provide very diverse results.

Appendix A

Real GDP % change YoY

Table 1					
Periods	IT GDP (REAL, %YOY)	GR GDP (REAL, %YOY)	IR GDP (REAL, %YOY)	PT GDP (REAL, %YOY)	ES GDP (REAL, %YOY)
1998	1.49	4.15	7.84	4.8	4.31
1999	1.52	3.05	11.03	3.9	4.48
2000	3.91	4.23	10.65	3.79	5.29
2001	1.61	3.61	5.28	1.94	4
2002	0.25	4	5.95	0.77	2.88
2003	0.24	5.81	3.69	-0.93	3.19
2004	1.37	4.78	6.8	1.81	3.17
2005	1.15	0.8	5.79	0.77	3.72
2006	2.1	5.58	5.88	1.55	4.17
2007 A	0.665	1.585	1.875	1.245	1.885
2007	1.33	3.17	3.75	2.49	3.77
2007 B	0.665	1.585	1.875	1.245	1.885
2008	-1.07	-0.23	-4.4	0.2	1.12
2009	-5.52	-4.31	-4.61	-2.98	-3.57
2010	1.65	-5.46	2	1.9	0.01
2011	0.72	-9.18	-0.06	-1.83	-1
2012	-2.85	-7.32	-1.08	-4.03	-2.93
2013	-1.75	-3.17	1.08	-1.13	-1.71
2014	0.17	0.39	8.45	0.89	1.38
2015	0.61	-0.31	26.29	1.6	3.2
2016	0.83	0.04	4.26	1.16	3.22
average of pre-crisis period	1.4305	3.7595	6.4785	1.9645	3.7095
average of post-crisis period	-0.6545	-2.7965	3.3805	-0.2975	0.1605
difference %	-1.4575	-1.7438	-0.4781	-1.1514	-0.9567

Appendix B

Share of total market capitalisation

<u>Greece</u>	
Sectors Average Market Value (%)	Percentages
GREECE-DS Financials - MARKET VALUE	46.7699
GREECE-DS Telecom - MARKET VALUE	15.7975
GREECE-DS Consumer Svs - MARKET VALUE	12.5933
GREECE-DS Consumer Gds - MARKET VALUE	7.3617
GREECE-DS Basic Mats - MARKET VALUE	6.4210
<u>Italy</u>	
Sectors Average Market Value (%)	Percentages
ITALY-DS Financials - MARKET VALUE	37.08441871
ITALY-DS Oil & Gas - MARKET VALUE	14.1916548
ITALY-DS Utilities - MARKET VALUE	12.30382934
ITALY-DS Telecom - MARKET VALUE	11.45146251
ITALY-DS Consumer Gds - MARKET VALUE	7.684180825
<u>Portugal</u>	
Sectors Average Market Value (%)	Percentages
PORTUGAL-DS Financials - MARKET VALUE	23.405312
PORTUGAL-DS Utilities - MARKET VALUE	16.45241574
PORTUGAL-DS Consumer Svs - MARKET VALUE	15.03660562
PORTUGAL-DS Telecom - MARKET VALUE	14.22134026
PORTUGAL-DS Industrials - MARKET VALUE	8.135696006

Appendix C

Critical values of the Dickey – Fuller GLS unit root test

Table 1. Dickey-Fuller GLS critical values

	1%	5%	10%
<i>Panel A – Intercept model</i>			
	-2.74	-1.96	-1.60
<i>Panel B – Intercept and trend model</i>			
50	-3.77	-3.19	-2.89
100	-3.58	-3.03	-2.74
200	-3.46	-2.93	-2.64
∞	-3.48	-2.89	-2.57

Appendix D

Proof of equation (8)

For $t=3$ and $q=2$, equation (7) is

$$V_3 = K_1 V_2 + K_2 V_1 + \Gamma X_3 + \varepsilon_3$$

And equation (8) is

$$\Delta V_3 = \Pi V_2 + \sum_{i=1}^1 A_i \Delta V_{3-i} + \Gamma X_3 + \varepsilon_3 = \left(\sum_{i=1}^2 K_i - I \right) V_2 + A_1 \Delta V_2 + \Gamma X_3 + \varepsilon_3$$

$$\Leftrightarrow \Delta V_3 = (K_1 + K_2 - I) V_2 + \left(- \sum_{j=2}^2 K_j \right) \Delta V_2 + \Gamma X_3 + \varepsilon_3$$

$$\Leftrightarrow \Delta V_3 = K_1 V_2 + K_2 V_2 - I V_2 - K_2 \Delta V_2 + \Gamma X_3 + \varepsilon_3$$

$$\Leftrightarrow V_3 - V_2 = K_1 V_2 + K_2 V_2 - V_2 - K_2 (V_2 - V_1) + \Gamma X_3 + \varepsilon_3$$

$$\Leftrightarrow V_3 = K_1 V_2 + K_2 V_2 - K_2 V_2 + K_2 V_1 + \Gamma X_3 + \varepsilon_3$$

$$\Leftrightarrow V_3 = K_1 V_2 + K_2 V_1 + \Gamma X_3 + \varepsilon_3$$

Both equation (7) and (8) give the same result, so these two equations are equal.

Appendix E

Serial correlations of the VAR models for the lag selection

Table A
Serial Correlation Results

Lags	<i>GREECE</i>		<i>ITALY</i>		<i>PORTUGAL</i>	
	Pre-crisis LM-Stat	Post-crisis LM-Stat	Pre-crisis LM-Stat	Post-crisis LM-Stat	Pre-crisis LM-Stat	Post-crisis LM-Stat
1	11.569	17.218	29.527	32.063	23.263	17.127
2	19.482	30.302	26.220	28.357	10.220	33.471
3	24.152	24.966	16.963	37.358*	19.927	22.915
4	30.596	29.030	21.378	27.069	15.179	35.773
5	22.922	27.227	22.827	20.645	10.776	32.739
6	22.278	26.582	17.311	14.938	20.575	20.429

Note: The null hypothesis is the no existence of serial correlation to the residuals of the VAR or VECM model. It is rejected based on the LM-statistics. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively. There are 16 degrees of freedom.

Appendix F

Greece

Appendix F.1 Impulse responses

Figure 1: Pre-crisis period

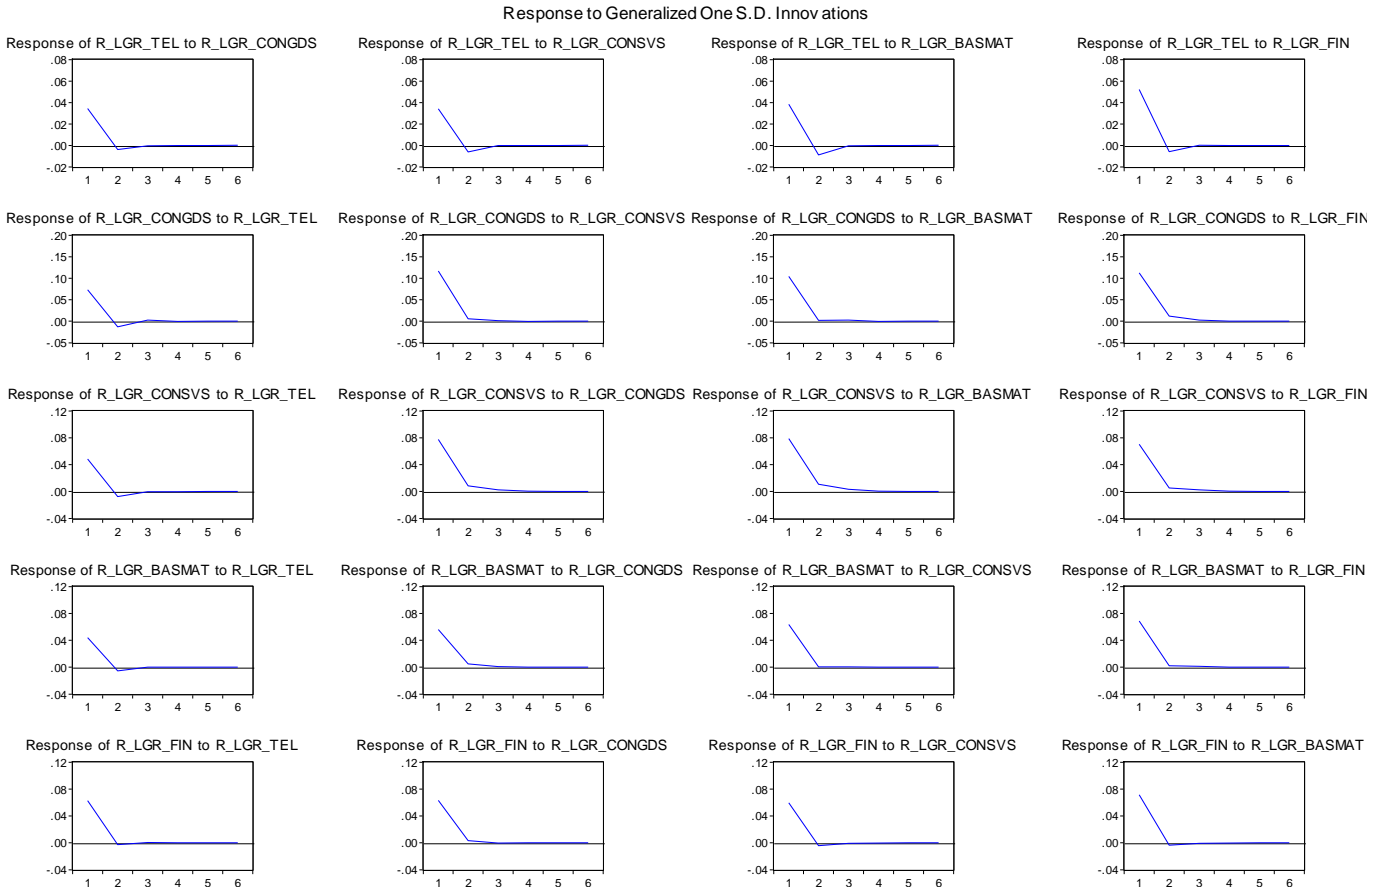
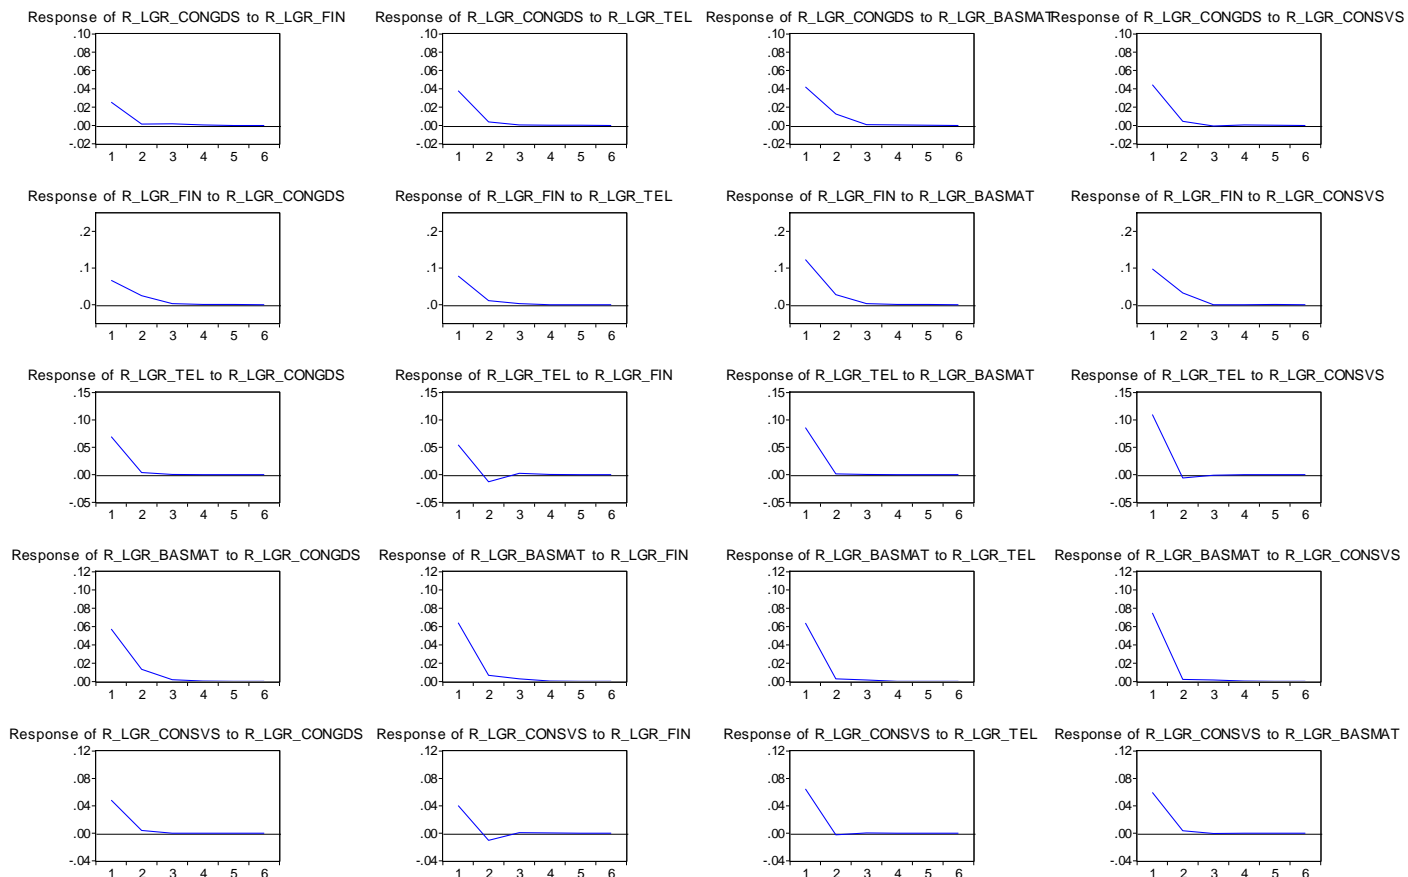


Figure 2: Post-crisis period

Response to Generalized One S.D. Innovations



Appendix F.2 Variance decomposition

Variance decomposition for the pre-crisis period - Greece

<i>Variance decomposition of R_LGR_TEL</i>					
Month	R_LGR_TEL	R_LGR_CONGDS	R_LGR_CONSVS	R_LGR_BASMAT	R_LGR_FIN
1	100.0000	0.0000	0.0000	0.0000	0.0000
2	99.0675	0.0463	0.2585	0.5800	0.0478
3	99.0542	0.0536	0.2586	0.5849	0.0486
4	99.0537	0.0537	0.2587	0.5851	0.0488
5	99.0537	0.0537	0.2587	0.5851	0.0488
<i>Variance decomposition of R_LGR_CONGDS</i>					
Month	R_LGR_TEL	R_LGR_CONGDS	R_LGR_CONSVS	R_LGR_BASMAT	R_LGR_FIN
1	18.8902	81.1098	0.0000	0.0000	0.0000
2	18.7501	78.4095	0.2171	0.0033	2.6201
3	18.7609	78.3895	0.2243	0.0065	2.6188
4	18.7608	78.3882	0.2244	0.0066	2.6200
5	18.7608	78.3881	0.2244	0.0066	2.6200
<i>Variance decomposition of R_LGR_CONSVS</i>					
Month	R_LGR_TEL	R_LGR_CONGDS	R_LGR_CONSVS	R_LGR_BASMAT	R_LGR_FIN
1	18.7186	31.2561	50.0253	0.0000	0.0000
2	18.5783	31.6432	49.3216	0.4329	0.0241
3	18.5539	31.6590	49.2738	0.4872	0.0262
4	18.5533	31.6597	49.2720	0.4886	0.0264
5	18.5533	31.6597	49.2719	0.4886	0.0265
<i>Variance decomposition of R_LGR_BASMAT</i>					
Month	R_LGR_TEL	R_LGR_CONGDS	R_LGR_CONSVS	R_LGR_BASMAT	R_LGR_FIN
1	23.6085	20.5594	11.3461	44.4860	0.0000
2	23.5937	20.9900	11.2131	44.1331	0.0701
3	23.5858	20.9892	11.2096	44.1249	0.0905
4	23.5857	20.9891	11.2101	44.1245	0.0906
5	23.5857	20.9891	11.2101	44.1245	0.0906
<i>Variance decomposition of R_LGR_FIN</i>					
Month	R_LGR_TEL	R_LGR_CONGDS	R_LGR_CONSVS	R_LGR_BASMAT	R_LGR_FIN
1	43.7944	17.6798	2.3035	8.3594	27.8629
2	43.2829	17.6997	2.9915	8.3086	27.7173
3	43.2618	17.6935	2.9951	8.3120	27.7376
4	43.2611	17.6933	2.9960	8.3126	27.7370
5	43.2611	17.6934	2.9960	8.3126	27.7370

Variance decomposition for the post-crisis period - Greece

Variance decomposition of R_LGR_CONGDS

Month	R_LGR_CONGDS	R_LGR_FIN	R_LGR_TEL	R_LGR_BASMAT	R_LGR_CONSVS
1	100.000	0.0000	0.0000	0.0000	0.0000
2	96.8874	0.0033	0.0446	2.8986	0.1661
3	96.7498	0.0352	0.0502	2.8999	0.2650
4	96.7446	0.0386	0.0503	2.9006	0.2659
5	96.7445	0.0386	0.0503	2.9007	0.2659

Variance decomposition of R_LGR_FIN

Month	R_LGR_CONGDS	R_LGR_FIN	R_LGR_TEL	R_LGR_BASMAT	R_LGR_CONSVS
1	10.0583	89.9417	0.0000	0.0000	0.0000
2	10.9684	86.7236	0.0018	0.9592	1.3470
3	10.9721	86.6491	0.0114	0.9814	1.3861
4	10.9717	86.6466	0.0122	0.9814	1.3881
5	10.9717	86.6466	0.0122	0.9814	1.3881

Variance decomposition of R_LGR_TEL

Month	R_LGR_CONGDS	R_LGR_FIN	R_LGR_TEL	R_LGR_BASMAT	R_LGR_CONSVS
1	22.3736	5.4273	72.1991	0.0000	0.0000
2	21.6968	6.2123	70.9503	1.1061	0.0345
3	21.6592	6.2289	70.8503	1.1181	0.1434
4	21.6583	6.2296	70.8487	1.1182	0.1453
5	21.6583	6.2295	70.8487	1.1182	0.1453

Variance decomposition of R_LGR_BASMAT

Month	R_LGR_CONGDS	R_LGR_FIN	R_LGR_TEL	R_LGR_BASMAT	R_LGR_CONSVS
1	27.4683	19.6092	7.4769	45.4455	0.0000
2	28.2525	19.2006	7.4882	44.5768	0.4819
3	28.2557	19.2233	7.4812	44.5497	0.4902
4	28.2567	19.2233	7.4811	44.5487	0.4902
5	28.2567	19.2233	7.4811	44.5487	0.4902

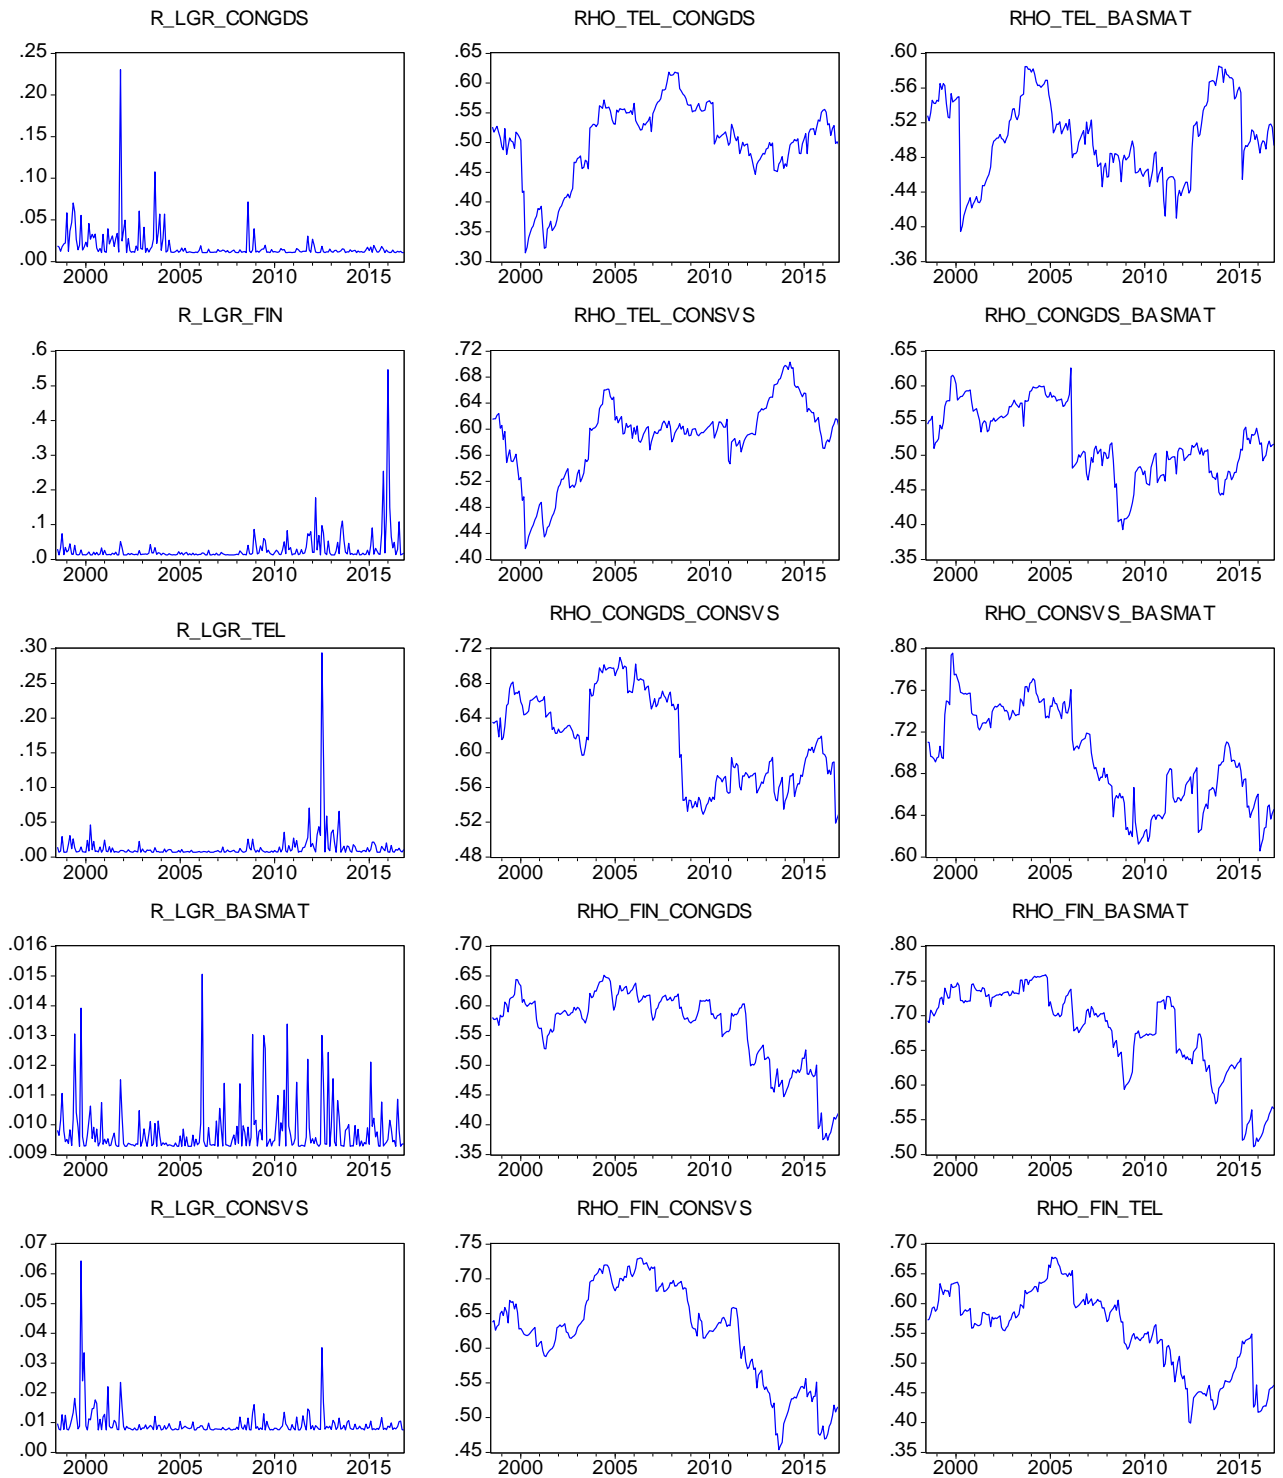
Variance decomposition of R_LGR_CONSVS

Month	R_LGR_CONGDS	R_LGR_FIN	R_LGR_TEL	R_LGR_BASMAT	R_LGR_CONSVS
1	30.6698	9.4068	23.4493	3.4122	33.0619
2	29.7372	10.9910	22.6093	4.8239	31.8386
3	29.6764	10.9827	22.5635	4.8289	31.9484
4	29.6747	10.9858	22.5632	4.8286	31.9477
5	29.6747	10.9858	22.5632	4.8286	31.9477

Appendix F.3 Time-varying volatilities and correlations

Time-varying volatilities

Time-varying correlations



Appendix G

Italy

Appendix G.1 Impulse responses

Figure 1: Pre-crisis period

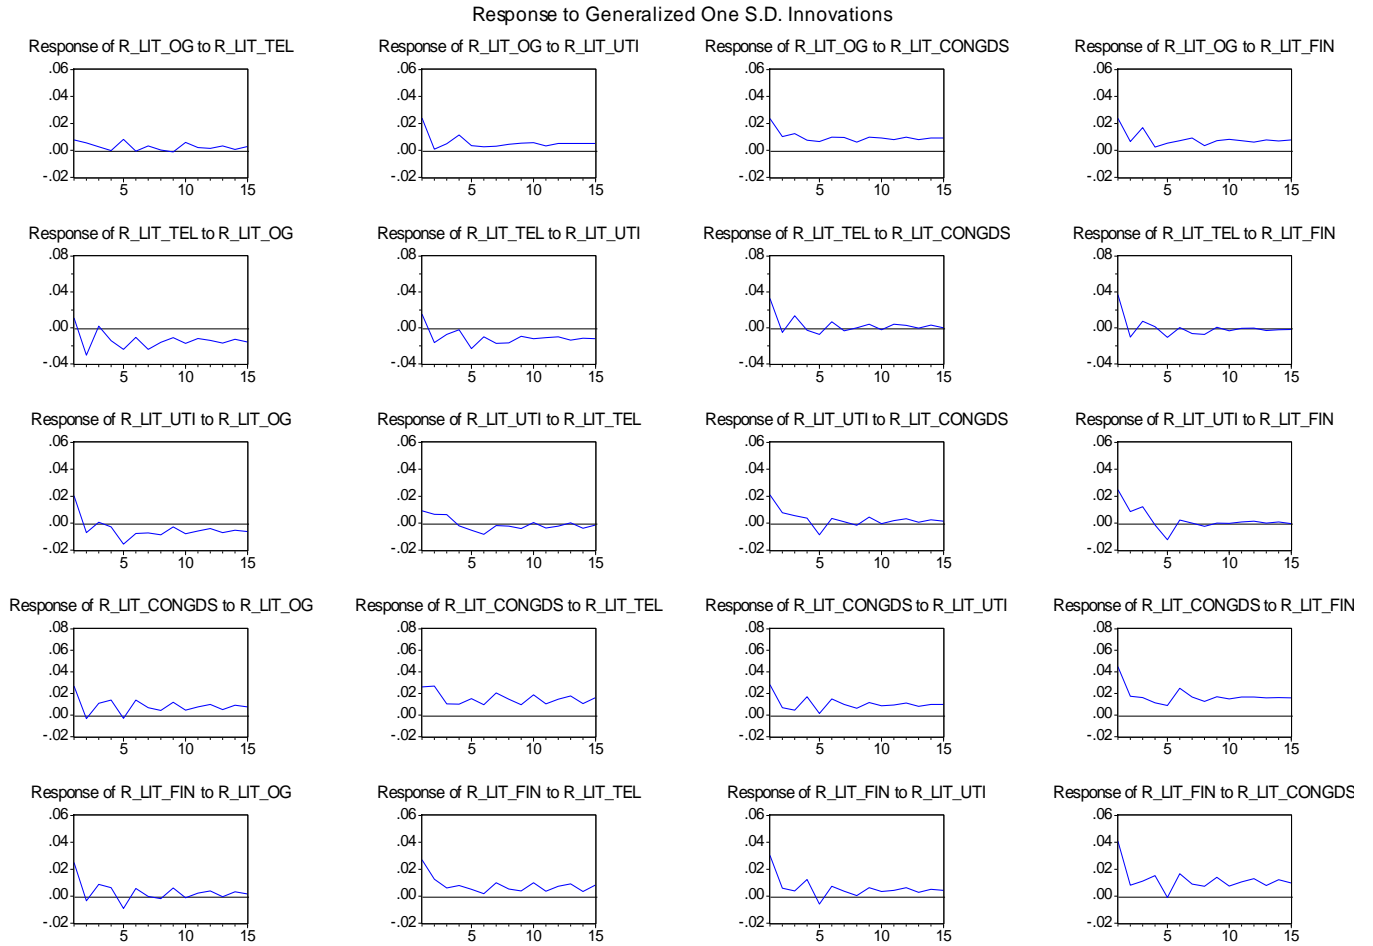
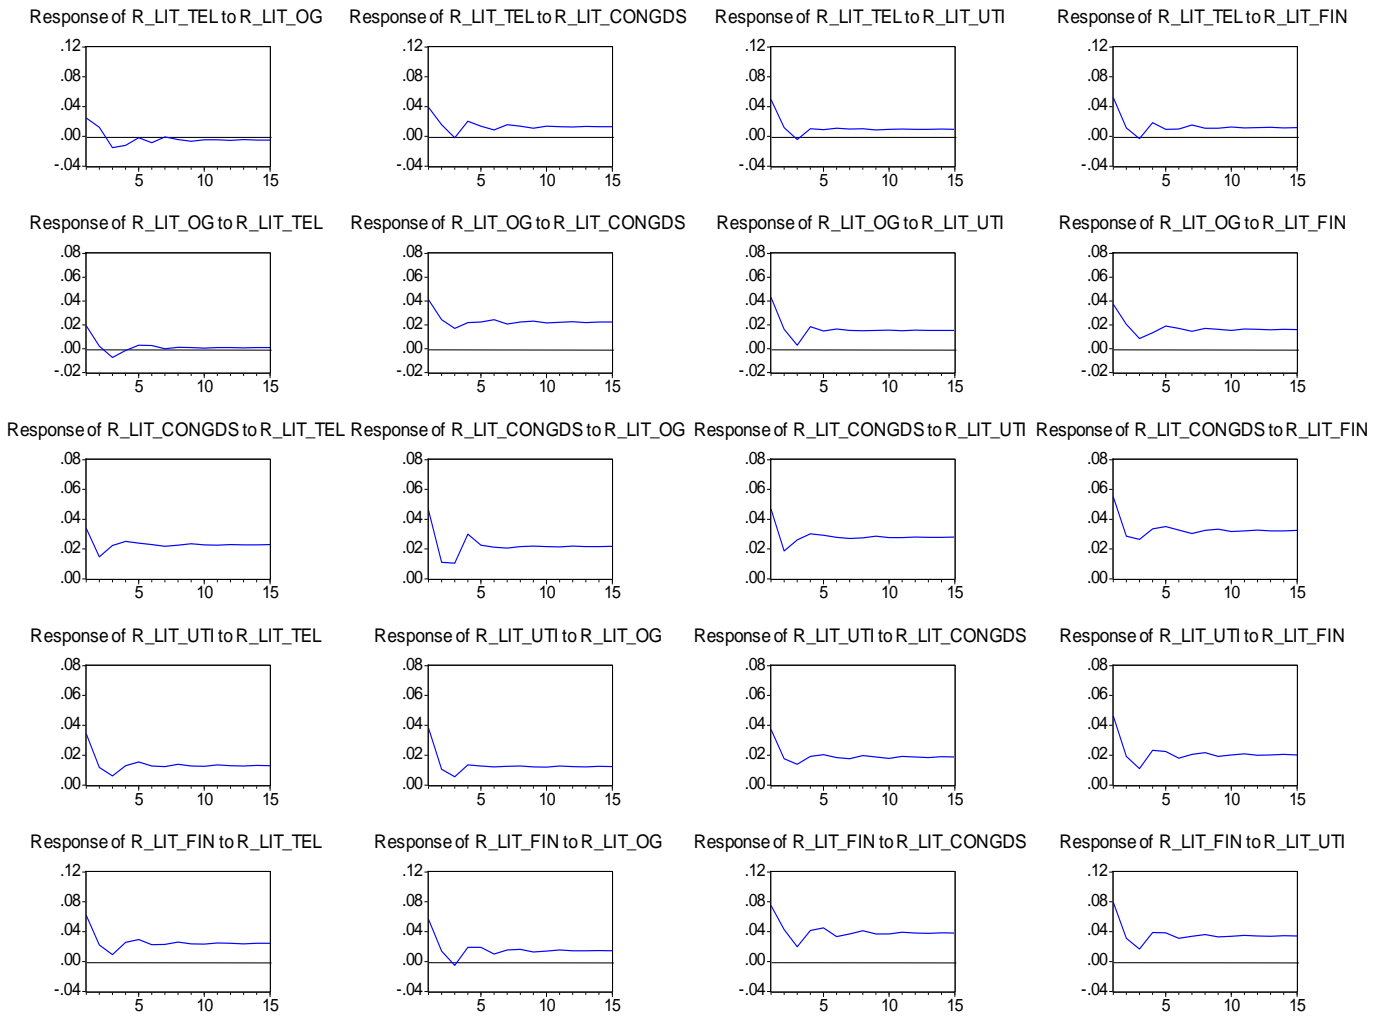


Figure 2: Post-crisis period

Response to Generalized One S.D. Innovations



Appendix G.2 Variance decomposition

Variance decomposition for the pre-crisis period - Italy

Variance decomposition of R_LGR_OG

Month	R_LGR_OG	R_LGR_TEL	R_LGR_UTI	R_LGR_CONGDS	R_LGR_FIN
1	100.000	0.0000	0.0000	0.0000	0.0000
2	96.9814	0.5996	0.5161	1.8599	0.0431
3	91.2305	0.5256	0.6395	3.2146	4.3898
4	89.0041	0.6098	1.5900	2.9688	5.8273
5	88.0935	1.6627	1.7036	2.8378	5.7024

Variance decomposition of R_LGR_TEL

Month	R_LGR_OG	R_LGR_TEL	R_LGR_UTI	R_LGR_CONGDS	R_LGR_FIN
1	2.0167	97.9833	0.0000	0.0000	0.0000
2	12.6748	84.5601	0.9112	0.1211	1.7328
3	11.7309	82.7252	2.5238	1.2713	1.7489
4	12.6409	81.7089	2.2931	1.7667	1.5904
5	16.9264	75.4497	4.1953	1.9632	1.4653

Variance decomposition of R_LGR_UTI

Month	R_LGR_OG	R_LGR_TEL	R_LGR_UTI	R_LGR_CONGDS	R_LGR_FIN
1	19.8193	1.9209	78.2599	0.0000	0.0000
2	19.3488	4.1106	74.3973	1.7249	0.4185
3	17.8622	5.2683	70.3607	1.6870	4.8218
4	17.4981	5.1689	69.0310	2.5839	5.7182
5	23.4910	4.9013	62.0716	2.5020	7.0341

Variance decomposition of R_LGR_CONGDS

Month	R_LGR_OG	R_LGR_TEL	R_LGR_UTI	R_LGR_CONGDS	R_LGR_FIN
1	19.2388	13.4759	5.6732	61.6122	0.0000
2	15.6691	27.0559	5.0981	52.0801	0.0968
3	16.5077	26.0454	4.6887	52.5212	0.2369
4	17.8355	24.1992	6.2104	49.4283	2.3266
5	17.0071	26.7966	5.8755	48.1217	2.1992

Variance decomposition of R_LGR_FIN

Month	R_LGR_OG	R_LGR_TEL	R_LGR_UTI	R_LGR_CONGDS	R_LGR_FIN
1	19.6761	17.7306	10.2597	14.8704	37.4632
2	18.3177	21.3154	10.4886	14.0836	35.7948
3	17.6535	18.9241	9.0343	13.4720	40.9161
4	17.3912	18.8329	10.6545	14.6860	38.4354
5	17.3026	17.7404	9.7208	13.2113	42.0249

Variance decomposition for the post-crisis period - Italy

Variance decomposition of R_LGR_TEL

Month	R_LGR_TEL	R_LGR_OG	R_LGR_CONGDS	R_LGR_UTI	R_LGR_FIN
1	100.0000	0.0000	0.0000	0.0000	0.0000
2	98.3258	0.8152	0.5611	0.0835	0.2144
3	94.3197	4.1412	1.1336	0.1522	0.2534
4	84.9222	7.5764	7.0986	0.1512	0.2516
5	83.5780	7.6964	8.0605	0.1429	0.5223

Variance decomposition of R_LGR_OG

Month	R_LGR_TEL	R_LGR_OG	R_LGR_CONGDS	R_LGR_UTI	R_LGR_FIN
1	7.970477	92.02952	0.000000	0.000000	0.000000
2	6.877233	82.13125	9.033846	1.059985	0.897684
3	7.023305	78.39723	11.85338	1.650987	1.075099
4	6.213941	73.53976	15.42038	3.545220	1.280699
5	5.726107	73.05402	16.24050	3.207115	1.772261

Variance decomposition of R_LGR_CONGDS

Month	R_LGR_TEL	R_LGR_OG	R_LGR_CONGDS	R_LGR_UTI	R_LGR_FIN
1	20.04719	25.32757	54.62524	0.000000	0.000000
2	18.09522	19.88106	61.60267	0.018299	0.402751
3	20.33218	16.59794	61.37526	1.163054	0.531563
4	22.08362	18.57353	57.91437	0.951165	0.477309
5	23.00643	17.79314	57.86430	0.888974	0.447155

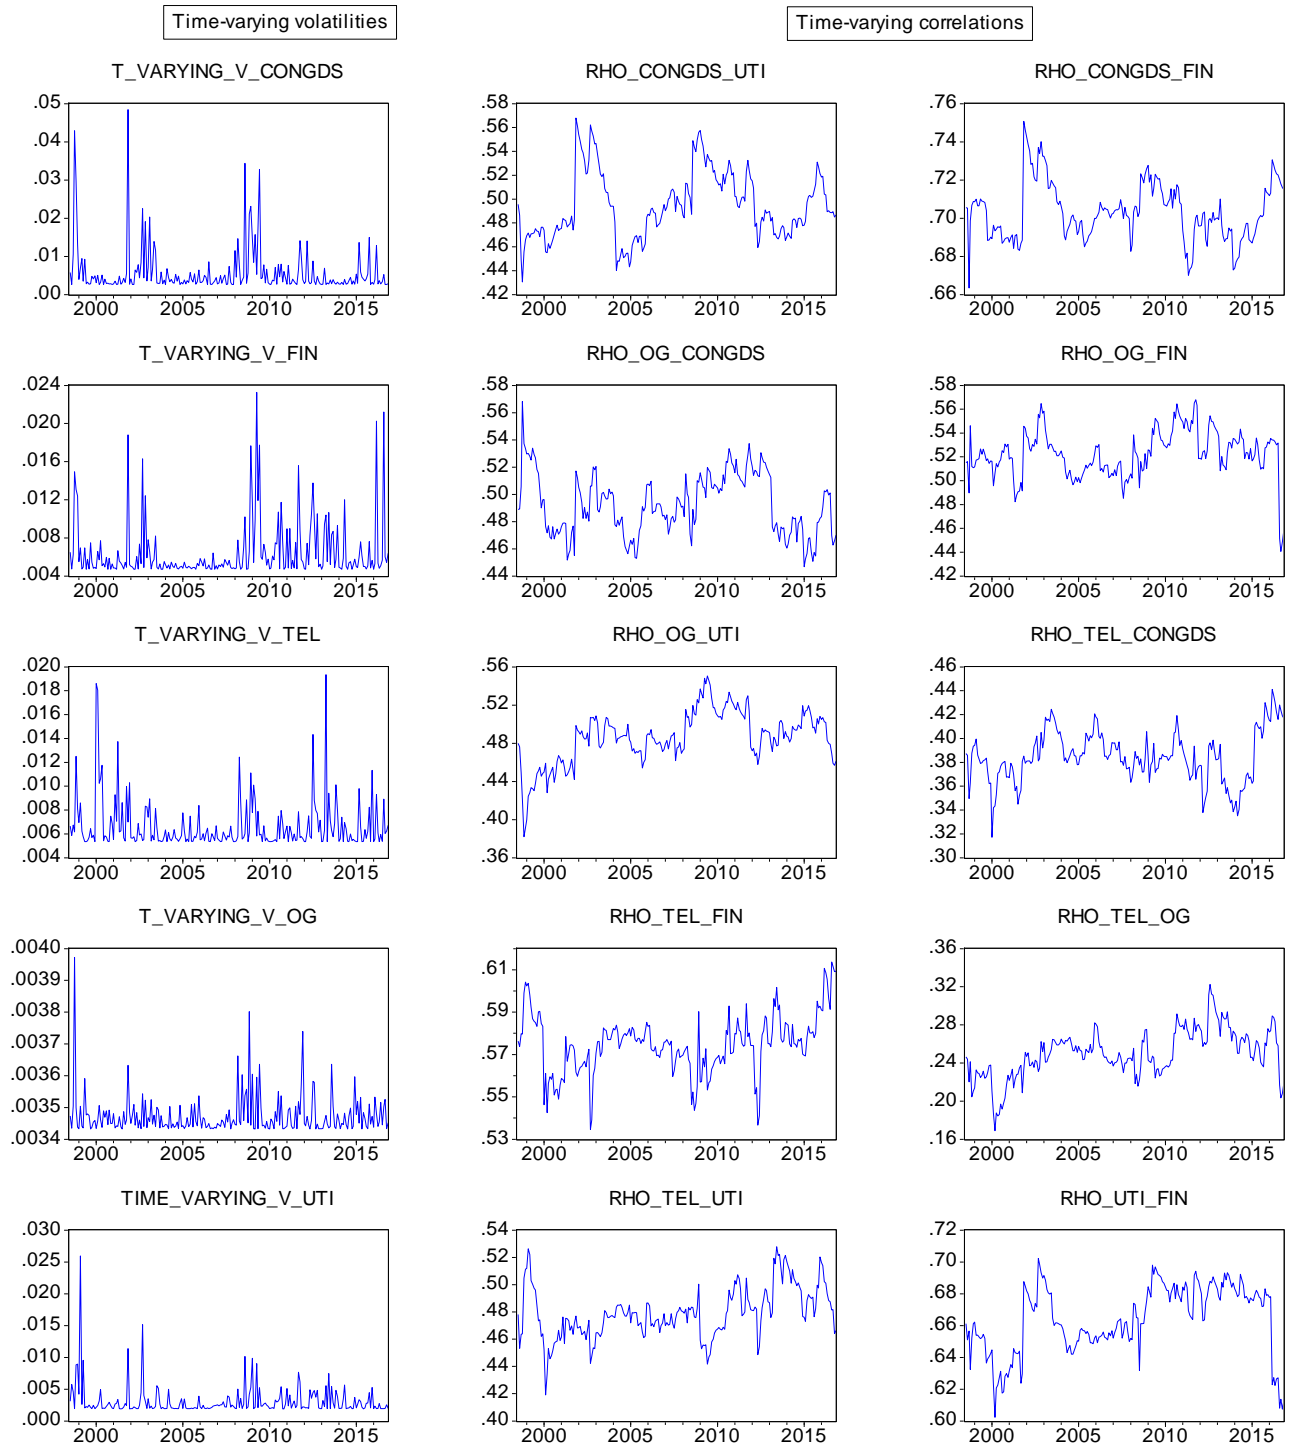
Variance decomposition of R_LGR_UTI

Month	R_LGR_TEL	R_LGR_OG	R_LGR_CONGDS	R_LGR_UTI	R_LGR_FIN
1	32.32411	24.55101	2.359355	40.76552	0.000000
2	32.02209	23.21679	5.344047	39.09540	0.321669
3	29.55555	21.16395	8.250158	40.11263	0.917710
4	27.04521	19.12110	8.931863	44.13485	0.766972
5	27.58911	18.05605	10.28541	43.35288	0.716547

Variance decomposition of R_LGR_FIN

Month	R_LGR_TEL	R_LGR_OG	R_LGR_CONGDS	R_LGR_UTI	R_LGR_FIN
1	35.42152	15.73558	12.29239	7.768815	28.78169
2	32.04247	13.10342	21.08854	7.541307	26.22427
3	29.99135	12.46480	24.11890	9.086740	24.33821
4	27.81977	10.88358	25.23507	9.672109	26.38947
5	27.85151	9.870564	27.41224	9.634174	25.23152

Appendix G.3 Time-varying volatilities and correlations



Appendix H

Portugal

Appendix H.1 Impulse responses

Figure 1: Pre-crisis period

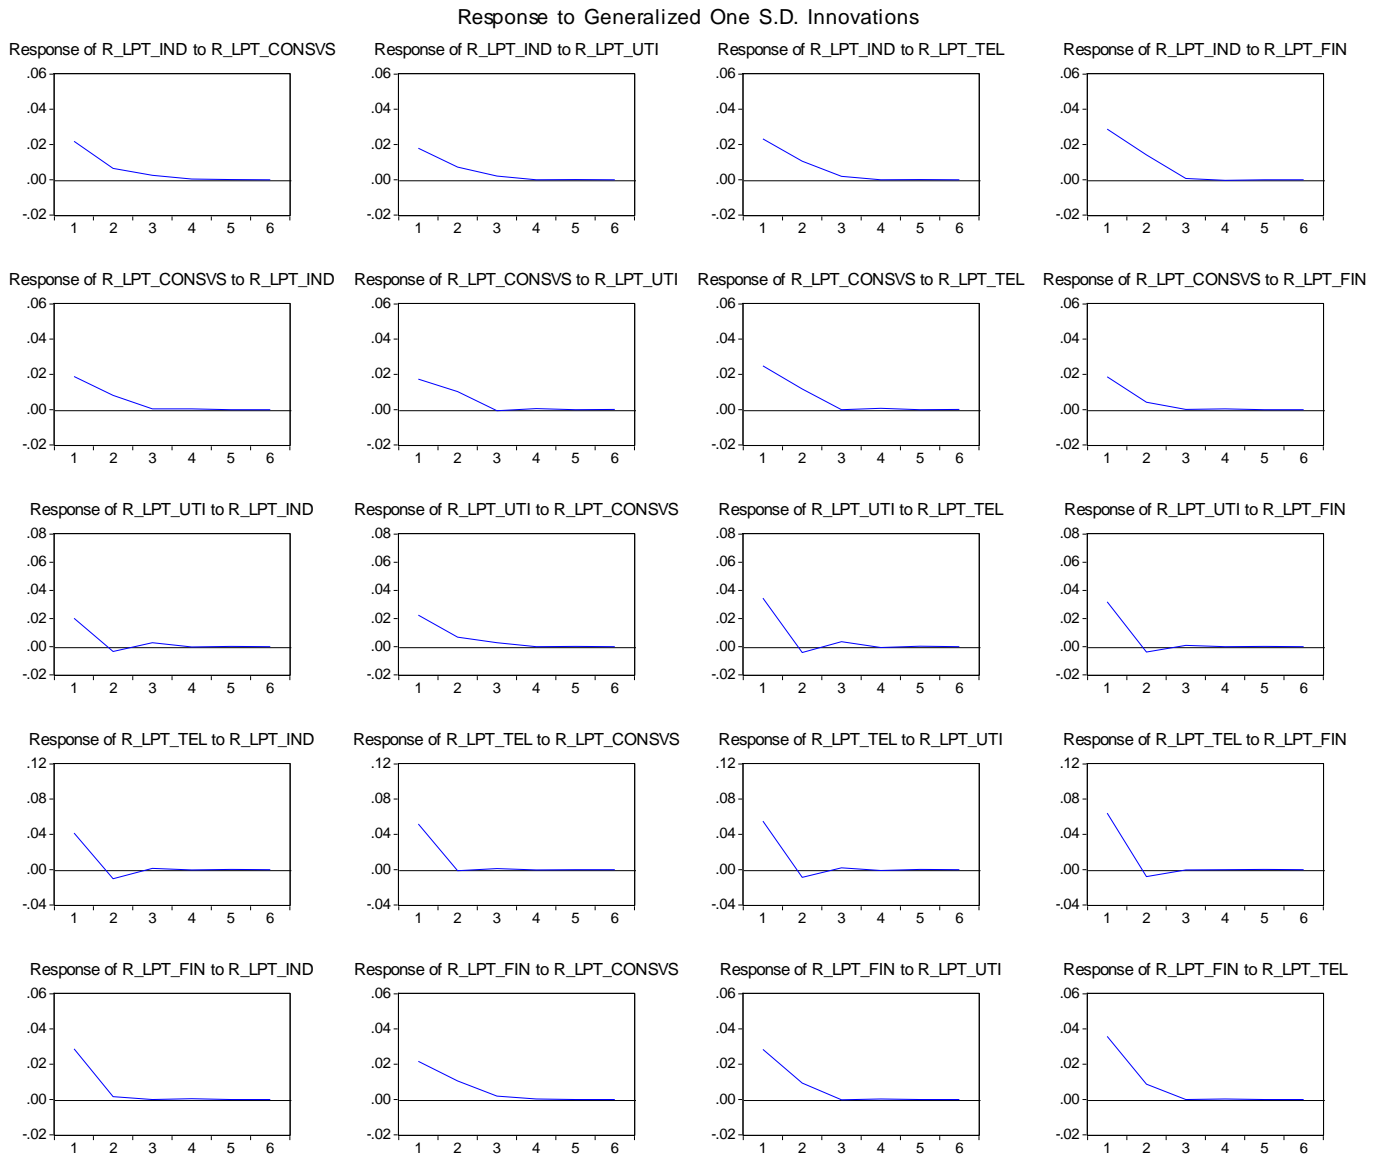
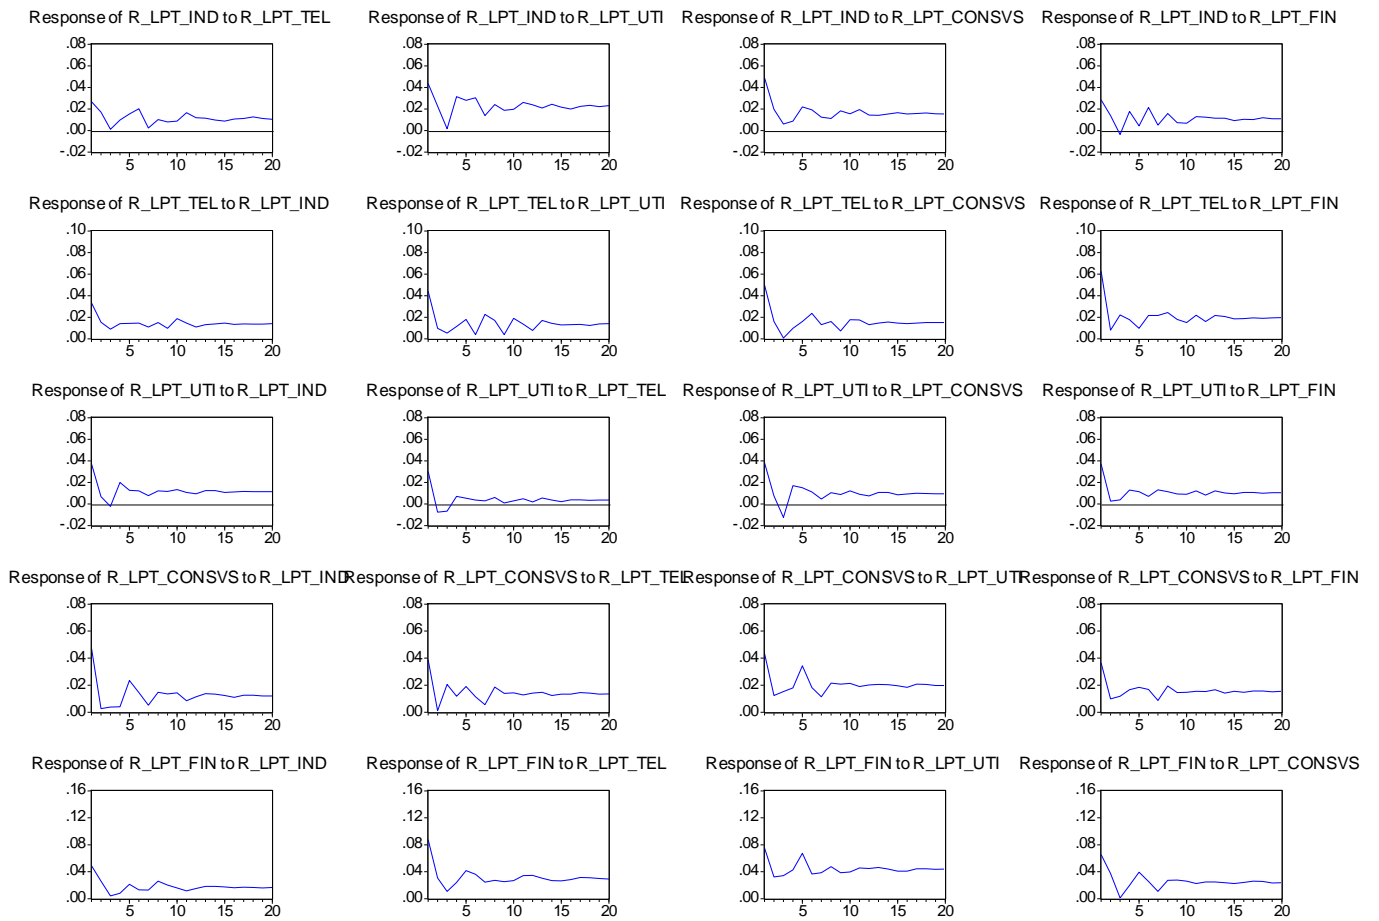


Figure 2: Post-crisis period

Response to Generalized One S.D. Innovations



Appendix H.2 Variance decomposition

Variance decomposition for the pre-crisis period – Portugal

Variance decomposition of R_LGR_IND

Month	R_LGR_IND	R_LGR_CONSVS	R_LGR_UTI	R_LGR_TEL	R_LGR_FIN
1	100.000	0.0000	0.0000	0.0000	0.0000
2	96.9814	0.5996	0.5161	1.8599	0.0431
3	91.2305	0.5256	0.6395	3.2146	4.3898
4	89.0041	0.6098	1.5900	2.9688	5.8273
5	88.0935	1.6627	1.7036	2.8378	5.7024

Variance decomposition of R_LGR_CONSVS

Month	R_LGR_IND	R_LGR_CONSVS	R_LGR_UTI	R_LGR_TEL	R_LGR_FIN
1	2.0167	97.9833	0.0000	0.0000	0.0000
2	12.6748	84.5601	0.9112	0.1211	1.7328
3	11.7309	82.7252	2.5238	1.2713	1.7489
4	12.6409	81.7089	2.2931	1.7667	1.5904
5	16.9264	75.4497	4.1953	1.9632	1.4653

Variance decomposition of R_LGR_UTI

Month	R_LGR_IND	R_LGR_CONSVS	R_LGR_UTI	R_LGR_TEL	R_LGR_FIN
1	19.8193	1.9209	78.2599	0.0000	0.0000
2	19.3488	4.1106	74.3973	1.7249	0.4185
3	17.8622	5.2683	70.3607	1.6870	4.8218
4	17.4981	5.1689	69.0310	2.5839	5.7182
5	23.4910	4.9013	62.0716	2.5020	7.0341

Variance decomposition of R_LGR_TEL

Month	R_LGR_IND	R_LGR_CONSVS	R_LGR_UTI	R_LGR_TEL	R_LGR_FIN
1	19.2388	13.4759	5.6732	61.6122	0.0000
2	15.6691	27.0559	5.0981	52.0801	0.0968
3	16.5077	26.0454	4.6887	52.5212	0.2369
4	17.8355	24.1992	6.2104	49.4283	2.3266
5	17.0071	26.7966	5.8755	48.1217	2.1992

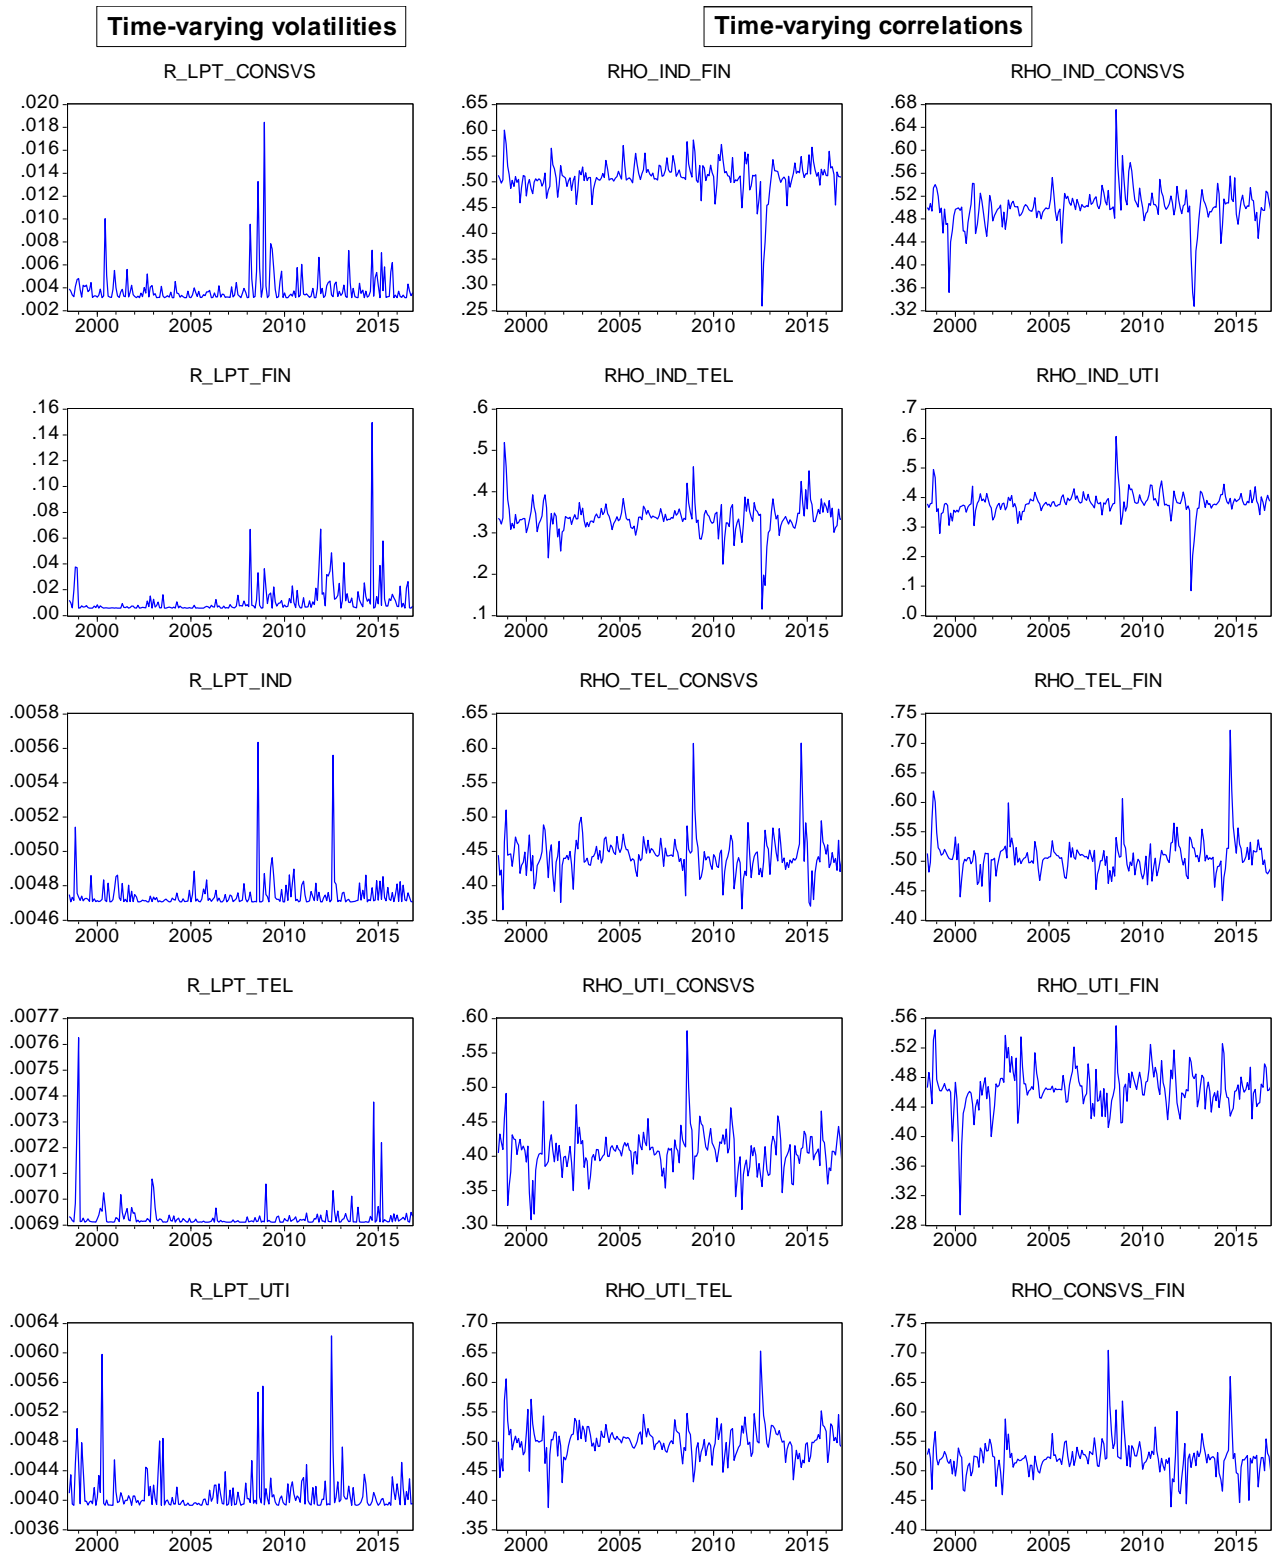
Variance decomposition of R_LGR_FIN

Month	R_LGR_IND	R_LGR_CONSVS	R_LGR_UTI	R_LGR_TEL	R_LGR_FIN
1	19.6761	17.7306	10.2597	14.8704	37.4632
2	18.3177	21.3154	10.4886	14.0836	35.7948
3	17.6535	18.9241	9.0343	13.4720	40.9161
4	17.3912	18.8329	10.6545	14.6860	38.4354
5	17.3026	17.7404	9.7208	13.2113	42.0249

Variance decomposition for the post-crisis period – Portugal

<i>Variance decomposition of R_LGR_IND</i>					
Month	R_LGR_IND	R_LGR_TEL	R_LGR_UTI	R_LGR_CONSVS	R_LGR_FIN
1	100.0000	0.0000	0.0000	0.0000	0.0000
2	91.1062	3.2062	4.2788	1.2245	0.1843
3	90.1368	3.2330	4.4854	1.2040	0.9409
4	79.5335	3.1493	14.1241	2.1522	1.0409
5	72.4827	3.8484	16.7095	2.3914	4.5680
<i>Variance decomposition of R_LGR_TEL</i>					
Month	R_LGR_IND	R_LGR_TEL	R_LGR_UTI	R_LGR_CONSVS	R_LGR_FIN
1	11.8563	88.1437	0.0000	0.0000	0.0000
2	13.9017	85.5817	0.0036	0.5018	0.0112
3	13.6519	80.5594	0.1466	1.7769	3.8652
4	14.8239	78.9761	0.1488	2.0004	4.0508
5	15.9540	77.0710	0.6774	1.9498	4.3478
<i>Variance decomposition of R_LGR_UTI</i>					
Month	R_LGR_IND	R_LGR_TEL	R_LGR_UTI	R_LGR_CONSVS	R_LGR_FIN
1	31.4009	7.8962	60.7029	0.0000	0.0000
2	30.1783	9.7168	56.5707	2.0410	1.4934
3	27.7023	9.6479	53.1029	6.3472	3.1997
4	28.5592	7.8084	55.8136	5.1465	2.6723
5	29.3707	7.4044	54.9504	5.5209	2.7535
<i>Variance decomposition of R_LGR_CONSVS</i>					
Month	R_LGR_IND	R_LGR_TEL	R_LGR_UTI	R_LGR_CONSVS	R_LGR_FIN
1	39.66951	10.12338	3.044973	47.16213	0.000000
2	36.82899	9.368486	5.967624	47.14492	0.689982
3	33.93295	14.82670	6.782853	43.44848	1.009020
4	31.73087	15.50745	9.925450	41.31087	1.525357
5	31.72194	14.01733	13.87323	38.88012	1.507375
<i>Variance decomposition of R_LGR_FIN</i>					
Month	R_LGR_IND	R_LGR_TEL	R_LGR_UTI	R_LGR_CONSVS	R_LGR_FIN
1	13.28433	31.03981	6.380882	0.254428	49.04054
2	14.66928	29.34216	6.449072	1.469725	48.06977
3	13.27900	26.84185	11.76319	2.465232	45.65072
4	12.15105	26.03841	16.93496	2.213535	42.66204
5	11.28502	25.30229	24.12125	1.904013	37.38743

Appendix H.3 Time-varying volatilities and correlations



References

- Ahmad, W., Sehgal, S. and Bhanumurthy, N.R. (2013). Eurozone crisis and BRIICKS stock markets: Contagion or market interdependence? *Economic Modelling*, 33, 209-225.
- Ahmed, W.M. (2012). On the interdependence structure of market sector indices: the case of Qatar Exchange. *Review of Accounting and Finance*, 11(4), 468-488.
- Andreoni, J. and Miller, J., (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2), 737-753.
- Arbeláez, H., Urrutia, J. and Abbas, N. (2001). Short-term and long-term linkages among the Colombian capital market indexes. *International Review of Financial Analysis*, 10(3), 237-273.
- Bartram, S.M. and Bodnar, G.M. (2009). No place to hide: The global crisis in equity markets in 2008/2009. *Journal of International Money and Finance*, 28(8), 1246-1292.
- Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Systems Research and Behavioral Science*, 9(3), 226-232.
- Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *Review of Economics and Statistics*, 72(3), 498-505.
- Brandimarte, P. (2006). *Numerical Methods in Finance and Economics*, John Wiley & Sons.
- Charness, G., Gneezy, U. and Imas, A., (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, 87, 43-51.
- Choi, S., Fisman, R., Gale, D. and Kariv, S., (2007). Consistency and heterogeneity of individual behavior under uncertainty. *American Economic Review*, 97(5), 1921-1938.
- Chouliaras, A.S., Christopoulos, A.G., Kenourgios, D. and Kalantonis, P. (2012). The PIIGS stock markets before and after the 2008 financial crisis: a dynamic cointegration and causality analysis. *International Journal of Banking, Accounting and Finance*, 4(3), 232-249.
- Dickey, D.A. and Fuller, W.A. (1979). Distribution of the Estimators for Autoregressive Time Series with a Unit Root. *Journal of the American Statistical Association*, 74, 427-431.
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.
- Elliott, G., Rothenberg, T.J. and Stock, J.H. (1996). Efficient Tests for an Autoregressive Unit Root. *Econometrica*, 64, 813-836.

- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339-350.
- Engle, R. F., & Ng, V. K. (1991). Time-varying volatility and the dynamic behavior of the term structure (No. w3682). *National Bureau of Economic Research*.
- Ewing, B.T., 2002. The transmission of shocks among S&P indexes. *Applied Financial Economics*, 12(4), 285-290.
- Frost, J. (2013). *Regression Analysis: How Do I Interpret R-squared and Assess the Goodness-of-Fit?* [Online]. Available at <http://blog.minitab.com/blog/adventures-in-statistics/regression-analysis-how-do-i-interpret-r-squared-and-assess-the-goodness-of-fit>
- Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3), 424-438.
- Griffiths, W.E., Hill, R.C., and Lim, G.C. (2008). *Using EViews for principles of econometrics*, John Wiley & Sons.
- Henderson V. and Hobson D. (2014). Risk aversion, indivisible timing options, and gambling, *Operations Research*, 61, 126-137.
- Henry, C. (1974). Investment decisions under uncertainty: the "irreversibility effect". *American Economic Review*, 64(6), 1006-1012.
- Hey, J. D., & Orme, C. (1994). Investigating generalisations of expected utility theory using experimental data. *Econometrica*, 62(6), 1291-1326.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644-1655.
- Hull, J. C. (2008). *Options, Futures and Other Derivatives*, Pearson Education International.
- Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica*, 59(6), 1551-1580.
- Kazemi, H. and Sohrabji, N. (2012). Contagion in Europe: examining the PIIGS crisis. *International Advances in Economic Research*, 18(4), 455-456.
- Koop, G., Pesaran, M.H. and Potter, S.M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1), 119-147.

- Levy, H., Levy, M., & Solomon, S. (2000). *Microscopic simulation of financial markets: from investor behavior to market phenomena*, Academic Press.
- Loomes G. and Pogrebna G. (2014). Measuring Risk Attitudes when Preferences are Imprecise, *Economic Journal*, 124(576), 569-593.
- Loomes, G. (1991). Evidence of a new violation of the independence axiom. *Journal of Risk and Uncertainty*, 4(1), 91-108.
- McDonald, R., & Siegel, D. (1986). The Value of Waiting to Invest. *The Quarterly Journal of Economics*, 101(4), 707-727.
- Musshoff O., Odening M., Schade C., Maart-Noelck S. C., and Sandri S. (2013). Inertia in disinvestment decisions: experimental evidence, *European Review of Agricultural Economics*, 40(3), 463-485.
- Oprea, R., Friedman, D., & Anderson, S. T. (2009). Learning to wait: A laboratory investigation. *The Review of Economic Studies*, 76(3), 1103-1124.
- Patra, T. and Poshakwale, S.S. (2008). Long-run and short-run relationship between the main stock indexes: evidence from the Athens stock exchange. *Applied Financial Economics*, 18(17), 1401-1410.
- Pesaran, H.H. and Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1), 17-29.
- Peskir, G., & Shiryaev A. (2006). *Optimal Stopping and Free-Boundary Problems*, Birkhäuser Verlag.
- Samitas, A. and Tsakalos, I. (2013). How can a small country affect the European economy? The Greek contagion phenomenon. *Journal of International Financial Markets, Institutions and Money*, 25, 18-32.
- Sandri S., Schade C., Mußhoff O., and Odening M. (2010). Holding on for too long? An experimental study on inertia in entrepreneurs' and non-entrepreneurs' disinvestment choices, *Journal of Economic Behavior and Organization*, 76, 30-44.
- Sheng, H.C. and Tu, A.H. (2000). A study of cointegration and variance decomposition among national equity indices before and during the period of the Asian financial crisis. *Journal of Multinational Financial Management*, 10(3), 345-365.

Tamakoshi, G. and Hamori, S. (2011). Transmission of stock prices amongst European countries before and during the Greek sovereign debt crisis. *Economics Bulletin*, 31(4), 3339-3353.

Tiwari, A.K., Mutascu, M.I. and Albuлесcu, C.T. (2016). Continuous wavelet transform and rolling correlation of European stock markets. *International Review of Economics & Finance*, 42, 237-256.

Vardhan, H., Sinha, P. and Vij, M. (2015). Behavior of Indian sectoral stock price indices in the post subprime crisis period. *Journal of Advances in Management Research*, 12(1), 15-29.

Wang, Z., Kutan, A.M. and Yang, J. (2005). Information flows within and across sectors in Chinese stock markets. *The Quarterly Review of Economics and Finance*, 45(4), 767-780.

Wilcox T. N. (2008). Stochastic Models for Binary Discrete Choice under Risk: A Critical Primer and Econometric Comparison, Risk Aversion in Experiments, *Research in Experimental Economics*, 12, 197–292.

Zhou W. and Hey J. D. (2016). Context Matters, *Working Paper, University of York*.